

SYNAPTIC SELF
HOW OUR BRAINS BECOME WHO WE ARE

BY
JOSEPH LeDOUX

ACKNOWLEDGMENTS

The bottom-line point of this book is “You are your synapses.” Synapses are the spaces between brain cells, but are much more. They are the channels of communication between brain cells, and the means by which most of what the brain does is accomplished.

To propose a synaptic explanation of the self, I had to discuss in some detail how the brain works. I’ve tried to do this without trivializing the facts— it’s not a pop psychology, how-to, or self-help book. Although writing about the brain in a way that will be clear to lay readers and at the same time not insulting to other scientists is tough, I’m pleased with the result.

But I didn’t achieve this goal alone. My wife, Nancy Princenthal, is an art critic, and a fantastic writer. She constantly urged me to use words economically, and to avoid repetition. She read and reread, each time (often to my annoyance) with a sharpened pencil in hand. And when I had finally met her critical approval, I turned the manuscript over to Rick Kot, my editor at Viking. It was great to have a pro like Rick on my side.

I also want to thank present and past members of my lab. Without their creativity and hard work, the science that inspired me to write *Synaptic Self* would not exist. And the work could not have been done without the generous support of the National Institute of Mental Health, New York University, the W H. Keck Foundation, and the Henry and Lucy Moses Fund.

I’m also grateful to colleagues who read certain chapters and commented on them. Thanks to Yadin Dudai, Jacek Debiec, Karim Nader, Greg Sullivan, and Rafi Lamprecht. In addition, several colleagues consulted with me on specific topics, including Tony Movshon, Carla Shatz, Dan Sanes, Jerome Kagan, Tim Wilson, Stephen Happel, Nancey Murphy, Barry Everitt, Norm White, Sandra File, Justine Kent, David Silbersweig, Jack Gorman, Amy Amsten, Liz Phelps, and Jonathan Cohen.

Will Chang went above and beyond on the Works Cited list. And Mian Hou and Claudia Farb made indispensable contributions to the illustrations.

Brett Kelly, Rick Kot’s assistant at Viking, was also very helpful throughout. Special thanks to Barbara Campo and Carla Bolte at Viking for their thoroughness and their patience with my corrections and to Don Homolka, eagle-eyed copy editor.

I also want to express my gratitude to my agents, Katinka Matson and [in Brockman. I greatly appreciate their advice and support.

Finally, I thank my children, Jacob and Milo. They inspire my synapses to change every day.

Writing a book is a humbling experience. You come to realize how many things you thought you understood but really didn’t, at least not well enough to explain them clearly. I learned a tremendous amount while writing *Synaptic Self* and hope you do too while reading it.

CHAPTER ONE

THE BIG ONE

DAD, WHAT IS THE MIND? IS IT JUST A SYSTEM OF IMPULSES OR SOMETHING TANGIBLE?

-- Bart Simpson

"I don't know, so maybe I'm not," the T-shirt said. Had it been another time or place, I probably would've chuckled and moved on. But I was on Bourbon Street when this postmodern homage to Descartes's proclamation, "I think, therefore I am," appeared on a young man weaving toward me. I had just had dinner in the French Quarter after a long day of data digestion and schmoozing at the Society of Neuroscience Conference, an annual get-together of thirty thousand or so brain researchers from around the world. The sound of Dixieland tunes, the aroma of stale beer, and the sight of scantily clad women dancing on runways inside dark bars had me contemplating the years and life changes that had come and gone since I myself had weaved down Bourbon Street during my college days in Louisiana. I headed back to my hotel, reflecting on my past and present life and wondering about what, all of a sudden, seemed to me to be the big question brain researchers should be asking:

"What makes us who we are?"

Neuroscience hasn't yet delved deeply into this puzzling issue. It has, for good reason, focused on how specific processes, like perception, memory, or emotion, work in the brain, but much less on how our brains make us who we are. I'd venture a guess that if a random sample of neuroscientists were

asked, "What do we know about the brain mechanisms of the self and personhood?" the predominant answer would be "Not much."

But maybe we know more than we think. Maybe some or even many pieces of the puzzle have already been discovered, and just have to be assembled into a coherent whole. Actually, I believe this might be the case. A lot of information is available about how the brain works, and while it may not yet be sufficient to fully explain persons, it should certainly encourage us to begin thinking about the problem.

My notion of personality is pretty simple: it's that your "self," the essence of who you are, reflects patterns of interconnectivity between neurons in your brain. Connections between neurons, known as synapses, are the main channels of information flow and storage in the brain. Most of what the brain does is accomplished by synaptic transmission between neurons, and by calling upon the information encoded by past transmission across synapses.

Given the importance of synaptic transmission in brain function, it should practically be a truism to say that the self is synaptic. What else could it be?

Not everyone, however, will be happy with this conclusion. Many will surely counter that the self is psychological, social, moral, aesthetic, or spiritual, rather than neural, in nature. My synaptic theory of the self is not proposed as an alternative to these views. It is, rather, an attempt to portray the way the psychological, social, moral, aesthetic, or spiritual self is realized.

I'll state unashamedly from the start that we can't, at this point, go all the way in formulating a complete

synaptic theory of personality. But even a partial understanding of the synaptic basis of who we are is, for me, an acceptable goal. For seeking knowledge about the brain is not only a valid scientific pursuit; it can also improve the quality of life, as when it uncovers new ways of treating neurological or psychiatric disorders.

SYNAPTIC NATURE

Let's start with a fact: People don't come preassembled, but are glued together by life. And each time one of us is constructed, a different result occurs. One reason for this is that we all start out with different sets of genes; another is that we have different experiences. What's interesting about this formulation is not that nature and nurture both contribute to who we are, but that they actually speak the same language. They both ultimately achieve their mental and behavioral effects by shaping the synaptic organization of the brain. The particular patterns of synaptic connections in an individual's brain, and the information encoded by these connections, are the keys to who that person is.

The genetic blueprint begins to unfold in the newly fertilized egg. Genes actually do two things in the broadest biological sense: they make us all the same (we're all humans), and they also distinguish us from one another (each of us has a unique genetic makeup that contributes to our individuality). When two people get together and make a baby, they always end up with a creature that looks and acts like a human and not like a monkey, dog, or fish. The common genetic heritage of our species dictates that the basic systems and molecules in my brain are the same as those in yours, and the basic mental and behavioral repertoire available to me is also available to you. We all walk upright, speak through our mouths, laugh, cry, and learn from experiences. But as children of a particular set of parents, who are themselves products of a particular genealogical history, we each also have genes that give our brains unique qualities and direct the specific manner in which the general mental and behavioral characteristics of our species are expressed. Genetic factors are in fact known to influence a variety of individual or personality characteristics, including how outgoing, fearful, or aggressive one is, as well as the likelihood that one will develop depression, anxiety, or schizophrenia.

Just how do genes affect individual behavior? In the simplest terms, they do so by making proteins that shape the way neurons get wired together. There are many steps, to say the least, between genes and their expression in actual behavior, but genetic sculpting of the synaptic organization of neural systems is the key to the process. Animal breeders have long known that it takes only a few generations of mating among carefully selected individuals to begin influencing behavioral traits of their offspring, such as how tame or fierce a given line of dogs is. When breeders try to customize behavior in this fashion, they in fact are often working with the synaptic organization of the brain. A few extra connections here, a little more or a little less neurotransmitter there, and animals begin to act differently. Once we realize that the basic wiring plan of the brain is under genetic influence, it's easy to see how not only animals but also people can have very similar brains and yet be so different, right from the start of their lives. Genetic forces, operating on the synaptic arrangement of the brain, constrain, at least to some extent, the way we act, think, and feel. The hopes of some, and the fears of others, surrounding the issue of human cloning obviously hinge on the importance of genes in regulating not only how we look, but also who we are.

For good reason, then, quite a lot of attention is being paid to genes these days. Widely read authors, including Richard Dawkins, E. O. Wilson, and Steven Pinker, have argued forcefully that key aspects of mind and behavior are inherited. Still, it's important to recognize that genes only shape the broad outline of mental and behavioral functions, accounting for at most 50 percent of a given trait, and in many instances for far less. Inheritance may bias us in certain directions, but many other factors dictate how one's genes are expressed.

For example, if a woman consumes excessive alcohol during pregnancy, or a child has a diet deficient in certain nutrients, a brain genetically destined for brilliance can instead turn out to be cognitively impaired. Likewise, a family history of extraversion can be squelched in an orphanage run with an iron fist, just as a natural tendency to be shy and withdrawn can be compensated for to some degree by the supportive encouragement of parents.¹ Even if it

becomes possible to clone a child who has died at a tender age, it's probable that the look-alike, having his own set of experiences, is going to act, think, and feel differently.

We hear a lot these days about how identical twins, reared apart by separate adoptive parents, can have similar habits and traits.' We hear less about the many ways they differ. The main outcome of Judith Rice Harris's controversial 1998 book, *The Nurture Assumption*, in which she proposed that parents hardly matter, was probably the emergence of clearer ideas about just how important, and under what conditions, parents *do* matter. The personality disorders of children raised in brutal Romanian orphanages are a shocking testimony to the fact that experiences can have profound effects on behavior.' Genes are important, but not all-important.

NURTURING NATURE

The puzzle of how nature and nurture shape who we are is simplified by the realization that synapses are the key to the operations of both. Whether your paycheck is deposited to your bank account automatically or you hand it over to the teller in person, it goes to the same place. Nature and nurture function similarly: they are simply two different ways of making deposits in the brain's synaptic ledgers.

For example, in life (and not just in cartoons) rodents really are afraid of cats. This is true of rats and mice living in the wild, but also of the rats studied in laboratories. Although lab rats come from breeding colonies that have been isolated from cats for many, many generations, any one of them, upon seeing a cat for the first time, will "freeze" dead in its tracks. This is Mother Nature at work, since the rat (and his immediate ancestors) never had the opportunity to learn from personal experience that cats are dangerous. This behavior is not just specific to rats or even animals, for people also freeze when faced with danger. Recall the amateur videotape of the bombing at the 1996 Atlanta Summer Olympics shown over and over on CNN. Within a second

•the explosion, everyone in the vicinity was crouching motionless."

Why did freezing become a genetically hardwired behavior? Obviously, there's no way to know for sure, since we have no fossil record tracing the evolution of behavioral or mental capacities. However, a reasonable hypothesis is that freezing is a beneficial response when faced with a predator. Predators, primal danger for most animals, respond to and are excited by movement. Freezing still in the face of danger is often the best thing for the prey to do. Because millions of years ago animals who did so were more likely to survive, today it's what most animals do, at least as an initial line of defense. Freezing not a choice but an automatic response, a preprogrammed way of dealing with danger. It sometimes backfires, however, as when a deer is frozen in the headlights of an oncoming car. Like most evolutionarily based strategies, it's good for many animals much of the time, but not for all animals all of the time.

What's interesting is that freezing also occurs if a rat (or other animal) clearly hears a sound that preceded an aversive stimulus (a mild electrical shock of its feet) on some prior occasion.' There's no predator around in this case, so how is the connection formed? The sound is a warning signal. Any rat that survives an encounter with a cat or other predator should store in its brain as much about the situation as possible so that the next time the sounds or smells that preceded the arrival of the cat occur, those stimuli can be attended to in order to increase its chances of staying alive. In the case of the electrical shock, pain receptors are activated and this substitutes for a close encounter of the harmful kind with a predator, making it possible for the stimuli that preceded the shock to be stored as if they were those that preceded a cat.

There are two basic ways in which these kinds of processes might have been wired up in the brain. First, a two-system operation might exist, with one system for responding to species-typical (innate or genetically programmed) dangers and another for learning about novel ones experienced by individuals in their lives. Or there might just be one system that takes care of both situations. In fact, experiments show that the latter is the way the brain actually works.

Damage to a region of the brain called the amygdala eliminates the tendency for rats to freeze to both the cat and to the sound. The amygdala is part of the brain system that controls freezing behavior and other defensive responses in threatening situations. Its synapses are wired by nature to respond to the cat, and by experience to

respond in the same way to dangers that are learned about. It's a wonderfully efficient way to do things:

rather than create a separate system to accommodate learning about new dangers, just enable the system that is already evolutionarily wired to detect danger to be modifiable by experience. The brain can, as a result, deal with novel dangers by taking advantage of evolutionarily fine-tuned ways of responding. All it has to do is create a synaptic substitution whereby the new stimulus can enter the circuits that the prewired ones used.

For the past twenty years, I've been trying to figure out how brains learn about dangers. In my previous book, *The Emotional Brain*, I described much of what is known about the organization of the brain system involved, and laid out the implications for understanding emotions. The basic wiring plan is simple: it involves the synaptic delivery of information about the outside world to the amygdala, and the control of responses that act back on the world by synaptic outputs of the amygdala. If the amygdala detects something dangerous via its inputs, then its outputs are engaged. The result is freezing, changes in blood pressure and heart rate, release of hormones, and lots of other responses that either are preprogrammed ways of dealing with danger or are aspects of body physiology that support defensive behaviors. Throughout my past research, I've made use of the fact that the fear system can learn and store information about stimuli that warn of impending bodily harm or other dangers. Recently, though, my lab has turned to studies of exactly how fear learning takes place at the synaptic level. This new work, described in this book, sheds light on how synaptic modifications in the fear system, particularly in the amygdala, allow us to benefit from past encounters with danger.

Most systems of the brain are plastic, that is, modifiable by experience. Which means that the synapses involved are changed by experience. But, as the fear example shows, learning is not the function that those systems originally were designed to perform. They were built instead to accomplish certain tasks (like detecting danger, finding food and mates, hearing sounds, or mov

ing a limb toward some desired object). Learning (synaptic plasticity) is just a feature that helps them do their job better.

Plasticity in all the brain's systems is an innately determined characteristic. This may sound like a nature-nurture contradiction, but it is not. An innate capacity for synapses to record and store information is what allows systems to encode experiences. If the synapses of a particular brain system cannot change, this system will not have the ability to be modified by experience and to maintain the modified state. As a result, the organism will not be able to learn and remember through the functioning of that system. All learning, in other words, depends on the operation of genetically programmed capacities to learn. Learning involves the nurturing of nature.

LEARNING WHO WE ARE

The study of learning and memory processes in the brain has advanced rapidly in the last few decades. This book is in large part based upon this new understanding of encoding and storage and their origin in synaptic function. Learning, and its synaptic result, memory, play major roles in gluing a coherent personality together as one goes through life. Without learning and memory processes, personality would be merely an empty impoverished expression of our genetic constitution. Learning allows us to transcend our genes, or, as the novelist Salman Rushdie said, "Life teaches us who we

Our genes may bias the way we act, but the systems responsible for much of what we do and how we do it are shaped by learning. Although a rat is by nature afraid of cats, it will live longer if it learns where in its particular world cats are most likely to be encountered and what kinds of sounds and smells are present when cats are nearby. A rat that has stored this information in its fear system is more worldly, and much better off, than one that has not. Something similar occurs in many if not most other brain systems: the information they encode and store today will contribute importantly to how they function tomorrow.

Our knowledge of who we are, of the way we think about ourselves, of what others think of us, and of how we typically act in certain situations is in large part learned through experience, and this information is accessible to us through memory. Without learning and memory, we wouldn't know if the person we are today jibes with the one we

were yesterday or the one we expect to be tomorrow. Without learning and memory, a person would have the bare-bones personality provided by genes, but wouldn't know much about it.

But, as we will see throughout this book, learning and memory also contribute to personality in ways that exceed explicit self-knowledge. The brain, in other words, learns and stores many things in networks that function outside of conscious awareness. These learned tendencies affect all aspects of mind and behavior, and are probably at least as important for day-to-day functioning as what we know about ourselves consciously.

THE BIG ONE, OR IS IT?

My walk down Bourbon Street got me thinking about the origins of the self as the key overarching problem that neuroscience should aim to solve. Many other brain scientists, however, might be inclined to say that consciousness is the big one.

Neuroscientists traditionally have avoided confronting consciousness. The topic was one that retired neuroscientists, facing their own mortality, would talk about, but young brain researchers knew better. Even joking about it could give you a bad reputation. But times have changed and discussions of consciousness by neuroscientists are on the rise. They have even been the basis of an indictment of science. In *The End of Science*, John Horgan, a social critic turned science journalist turned science critic, announced that neuroscience was, like other sciences, dead. But unlike some other disciplines, he argued, neuroscience was ending not because it had solved its big question—how does consciousness work?—but because it never would.

I think it's good that scientists are now interested in consciousness, but I also believe it's being overemphasized. Suppose next week we find out that after decades of false starts and failed promises, an indefatigable neuroscientist finally has solved the consciousness problem. Would that really tell us what makes people tick? Would we now understand why schizophrenia emerges in one person but not his twin brother; why, when two people are faced with bodily harm, one is paralyzed by his fear while the other fights back; why an excessively shy child is likely to become an anxious adult; why some people are vegetarians and others enjoy red meat; why a meat-eater sometimes orders a vegetarian dish; why my kids can't stand the music I listen to; or why I like theirs? The answer to each of these is clearly No!

The question *Synaptic Self* asks is not "How does consciousness come out of the brain?" but rather "How does our brain make us who we are?"

What a person is, and what he or she thinks, feels, and does, is by no stretch of the imagination influenced only by consciousness. Many of our

thoughts, feelings, and actions take place automatically, with consciousness only coming to know them as they happen, if at all. Figuring out the mechanism of consciousness would surely be a major scientific coup, but it wouldn't explain how the brain works, or how our brains make us the individuals we are.

An understanding of the mystery of personality crucially depends on figuring out the unconscious functions of the brain. *Unconscious* is actually a notoriously ambiguous word. Some people may think it refers to Freudian repressed memories, others that it is what happens in coma, or after being hit on the head, or drinking too much. None of these definitions is what I have in mind. Instead, what I mean by the term is the many things that the brain does that are not available to consciousness. (We can be thankful this is the way we function, for if we had to consciously plan each muscle contraction, our brain would be so busy we would probably never end up actually taking a step or uttering a sentence.)

Consciousness, at least the kind of consciousness we mean when we talk about our own mental states, very likely developed in the brain recently in evolutionary history, layered on top of all the other processes that already existed. Unconscious operation of the brain is thus the rule rather than the exception throughout the evolutionary history of the animal kingdom." It's a linguistic quirk, or a revealing cultural assumption, that the older (unconscious) processes are defined as negations of the newer one (consciousness). Language isn't perfect.

What, then, are all these unconscious processes? Actually, they include almost everything the brain does,

from standard body maintenance like regulating heart rate, breathing rhythm, stomach contractions, and posture, to controlling many aspects of seeing, smelling, behaving, feeling, speaking, thinking, evaluating, judging, believing, and imagining² We can be and often are aware of what we are doing when these things happen, but much of the time consciousness is informed after the fact. When someone speaks to you, for example, you decode sentence meaning on the basis of the sound of the words (phonology), the meaning of the words (semantics), the grammatical relations between the words (syntax), and your knowledge about the world (pragmatics). You usually are not aware of performing these operations, but simply do them. While you end up consciously knowing what the person said, you don't have conscious access to the processes that allowed you to comprehend the sentence. Similarly, when you yourself utter sentences, you go through the same processes, often without a conscious thought, but this time as the generator rather than the receiver. Our abilities to perceive the world, attend to objects and events, remember, imagine, and think all operate pretty much in this fashion. Collectively, these processes have been called the psychological or cognitive unconscious and they account for much of mental life.

Does that mean we'll know what a person is when we figure out the operations of these various conscious and unconscious functions? While unlocking the synaptic mechanisms underlying each of these processes is itself going to be quite a challenge, we need to go beyond the mere explanation of how each works in isolation. We need to understand how the many processes interact, and how the particular interactions that take place inside an individual's brain give rise to and maintain who he or she is. My aim in this book is to show how it is possible, at least in principle, to begin to understand the self in terms of such synaptic interactions.

Previous attempts to relate the self to the brain have mostly done so via the conscious self. Recently, though, it appears the tide is beginning to turn. Antonio Damasio's book, *The Feeling of What Happens*, discusses the *protoself* a kind of core self that exists outside of consciousness, and in *The Minds- Past* Michael Gazzaniga emphasizes the importance of unconscious processes in the production of consciousness. Thus, consciousness, after long being neglected in brain science, is finally getting the attention it deserves, but it is also being put in its place—as part of, but not comprising the entirety of, the mental terrain. Although these other books deal with the self and the brain, they do not explore the biological mechanisms by which the brain makes the self. This is what *Synaptic Self* does.

As we begin to understand ourselves in synaptic terms, we don't have to sacrifice other ways of understanding existence. The idea that the self is created and maintained by arrangements of synaptic connections, in other words, doesn't diminish who we are. It instead provides a simple and plausible explanation for how the enormously complex psycho-spiritual-sociocultural package of protoplasm we call our self is possible.

CHAPTER TWO

SEEKING THE SELF

KNOW THYSELF.

—Oracle at Delphi

SELF-KNOWLEDGE IS A DANGEROUS THING.

—Loj Reeo

Before we go looking for the essence of a person in the brain, it would help to have some conception of what we are seeking. There's certainly no shortage of opinion about what terms like *personality* or *the self* might mean. William James, for example, proposed: "*In its widest possible sense... a man's-Self is the sum total of all that he CAN call his*, not only his body and his psychic powers, but his clothes and his house, his wife and children, his ancestors and friends, his reputation and works, his lands and horses, and yacht and bank-account.

If they wax and prosper, he feels triumphant; if they dwindle and die away, he feels cast down,—not necessarily in the same degree for each thing, but in much the same way for all."

In fact, a whole area of psychology is devoted to the study of personality and the self. And theologians, philosophers, novelists, and poets have also had much to say on the subject. While seemingly deep truths have sometimes emerged from these musings, it is not clear how, if at all, these insights might relate to the workings of the brain. After all, for many people, the brain and the self are quite different. I hope to show here that this is not the case. To do so, I first need to describe a way of thinking about the self that is compatible with current understanding of brain function.

"SOUL SERENADE"

A few months after starting this book, I attended a conference on the relation between the brain and the soul, sponsored by (fittingly enough) the Vatican. The specific topic was "Neuroscience and Divine Action," and the theologians who organized this meeting were trying to reconceptualize Church teachings in a way that would make sense in light of current scientific understanding of how the world works. In particular, they were attempting to determine how it is possible for God to influence people's lives without violating the laws of physics. I can't present the full range of views expressed, but one that stood out was the notion that God *interacts* but doesn't *intervene*.

The basic idea goes something like this: In the beginning, God set the universe up in a certain way (that is, he created the laws of physics) and has subsequently left it alone, at least for the most part. On a typical day, therefore, God doesn't control the position of stars and planets, move mountains, part seas, change weather conditions, or make people do things they wouldn't otherwise do (in other words, he doesn't usually intervene), but he does communicate with people (he interacts).

Our concern here is not with the theological arguments for and against a noninterventionist view of God but rather with the possibility (or impossibility) of a scientific view of interaction. Given that people live a physical existence in the physical world, and that God is not part of the physical world, the question is: How can God interact with people? If you believe in the existence of a nonmaterial soul, then all you need to assume is that when God was creating the universe, he worked out some way of interacting with the soul. Since both God and the soul are nonmaterial, that interaction would also be nonmaterial, and the laws of physics would therefore be unviolated when

interactions occur.

Much to my surprise, however, many of the theologians attending this meeting didn't believe in a classic nonmaterial soul (this would probably be an even bigger surprise to the faithful they represent). Instead, they seemed to accept the principle that the mind is inexorably tied to the brain, and they consequently believed in a soul that is pretty much one and the same as the neurally mediated mind, a part of the physical world that must by its nature obey the laws of physics.

Theologians who link the soul to the physical world actually have history on their side. While many Christians today continue to believe in a soul that is separate from the body and that survives death, this idea didn't really be-

come prominent in Christianity until the Middle Ages. Early Christian teachings emphasized the resurrection of the body itself on Judgment Day rather than merely the survival of a nonmaterial soul. According to Mark 9:47, Jesus said, "It is better for you to get into the Kingdom of God with but one eye than to be thrown into Gehenna with both eyes." This lesson was apparently not meant symbolically but instead reflected the early Judaic notion that we take into the afterlife the bodily state with which we leave this life. (This explains why Jews and Christians, and Muslims, have cemeteries on the western slopes of Mount Olive, facing the Eastern Gate of the Old City in Jerusalem: the closer their bodies are buried to the Eastern Gate, which is where the Final Judgment is expected, the sooner they will be raised.) Ancient Egyptians also seem to have believed that the body, and not just the spirit, carries on in some way after death. In the Tomb of Menena, for example, a foe of the royal family chiseled off their faces to ensure that they went into the next life physically challenged.

If the soul is indeed physical in nature, part of the dilemma about how to sustain a belief in both physics and God would be solved (the part about how the soul meshes with the body). However, the thoroughly modern theologians would still be in a bit of a quantum pickle. If the soul is equivalent to the mind, and the mind depends on the functioning of the brain, how can God interact with people without physically affecting their neurons and, thus, intervening? And where is the soul hanging out while the body decays in the interim between death and Judgment Day? Not surprisingly, the Vatican conference ended inconclusively. No matter how all the pieces of the puzzle were moved around, they didn't fit together to make a coherent picture. As the philosopher David Hume said long ago, logic and reasoning (and presumably science) cannot explain the immortality of the soul. Either you believe or you don't.

My reason for discussing this conference and the issues it raised is not so much to argue the point that it would be difficult, and maybe impossible, to find scientific solutions to theological riddles, but rather to demonstrate that a spiritual view of the self isn't (or doesn't have to be) *completely* incompatible with a biological one. Whatever else we are and aren't, much of what we are is accounted for by what goes on in our brains. Some theologians, as we've seen, have come to accept this. But even people who believe in an immaterial soul that survives death have acknowledged the fact that the normal functioning of the soul depends on the brain. Shakespeare embraced this notion when he called the brain the soul's frail dwelling. A few minutes with my mother, a devout Catholic with Alzheimer's disease, makes it painfully clear just how fragile the soul's dwelling is.

THE LIMITS OF LOGIC

Theologians aren't the only ones who have concerned themselves with interactions between body and soul. The question is also one of the major puzzles that has occupied philosophers through the ages. In the seventeenth century, the French mathematician René Descartes devised a way of thinking about body and soul that has shaped philosophical debate on this topic ever since.' Like the contemporary theologians described above, Descartes sought a means of reconciling science and faith. His solution was to propose that "the mental" and "the physical" were separate substances that met and interacted at a special place in the brain.

Historians seem to agree that the earliest Greek philosophers did not have clearly distinguished notions of body and soul." Later, though, some philosophers came to view body and soul as separate. Plato, for example, believed that the intellectual essence of an individual—his psyche or soul—survives death.' In fact, Plato looked forward to death so that he could be free of his body and all its needs and passions and finally be capable of pure

thought.’ For Aristotle, in contrast, body and soul were so integrally related that they could not be separated, though they could be distinguished conceptually.’ By the Middle Ages, philosophers adopted a combination of these notions, viewing body and soul as two unified ‘substances’ (like Aristotle) but regarding the soul as eternal (like Plato). Aquinas, for example, believed that the intellectual, nonmaterial qualities of mind gave the soul immortality; and that the body was resurrected and reunited with the soul on Judgment Day.’

This was the intellectual backdrop against which Descartes played his influential discussion of body and mind. Like Plato, he viewed the mental and physical as separate substances: “My soul, by which I am what I am ... is entirely and absolutely distinct from my body, and can exist without it.” In a fusion of faith and psychological theory, Descartes equated the soul with consciousness, and said only humans have conscious control of their behavior. Therefore, only human souls can gain or lose access to heaven by their actions. The behaviors of other animals were, in Descartes’s scheme, reflexive or automatic, and carried out without thought. So, for Descartes, if it wasn’t conscious, it wasn’t mental. Descartes didn’t exactly deny the existence of unconscious processes, but simply relegated them to the physical world, proposing that they function in humans the way they do in mindless (soulless) animals.

But if ‘the physical’ and ‘the mental’ are completely different entities, how can the conscious soul (the mental) be responsible for the physical body? Descartes’s solution was that the conscious soul substance can interact with the material body by means of a small region of the brain called the pineal gland. ‘While most parts of the brain exist in duplicate, the pineal gland is singular and centrally located, which suggested to Descartes that it must be the seat of mind-body interaction—a place where commands from the soul can influence the body, and where information from the body (either about the outside world or the body itself) can enter the soul as perceptions, emotions, and knowledge.

In Descartes’s scheme, the nonphysical soul substance actually served the dual function of communicating with the physical body as well as with God. This interplay between physical and nonphysical substances is precisely the kind of solution to the problem of how God interacts with people that the theologians at the Vatican conference were trying to move beyond. Our interest here, however, is not in the theological question (How does God interact with the soul?), but in the philosophical one (How does the mind interact with the body?). Descartes’s framing of the philosophical question, and his particular answer to it (a mind-body interaction in the brain), set up the conundrum known as the mind-body problem, which philosophers have struggled with ever since. I want to make two points about this subject that are relevant to the present discussion.

First, in equating the mind with consciousness, Descartes framed the mind-body problem in terms of the relation between consciousness and the brain. As such, the mind-body problem, in its traditional conception, concerns an *aspect* of the mind rather than the whole mind. Most of what the brain does is not, in fact, part of the traditional mind-body debate. Some contemporary philosophers take a broader view than Descartes and accept that certain nonconscious aspects of brain function also contribute to mental life.’ However, they regard such nonconscious aspects as “easy problems” that are not really the concern of philosophy. As will become apparent later in this chapter, I believe that these implicit or unconscious aspects of the self also play an important role, in fact an essential one, in shaping who we are and explaining why we do what we do.

Second, it’s important to distinguish the philosophical mind-body problem from the neuroscientific problem about how the brain creates the mind. Philosophers, by definition, seek philosophical solutions to problems (including the mind-body problem) and discuss the possible relations, the logic, that might exist between fundamental substances in nature (matter and mind). Neuroscientists, by contrast, typically start with the assumption that the materialist view of the mind-body problem is correct (that the mind is a product of the brain), and then try to understand how the brain makes the mind possible.’ In fact, many philosophers today accept some version of materialism, but even if the tides should shift in the coming years and dualism (the belief that mind and body are separate substances) should take over philosophy, neuroscientists will not be out of jobs. Brain researchers are, after all, studying the brain, not philosophy. This does not mean that the paths of the neuroscientist and philosopher never

cross. They often do, and when they do, each group has sometimes been enlightened (and sometimes enraged) by the other. But, ultimately, because philosophers and brain scientists are pursuing different concerns, progress in one field does not necessarily signal an advance or defeat in the other.

In spite of my own contention that consciousness is not the be-all and end-all of mind and behavior, I nevertheless have considerable sympathy for the belief that neuroscience will come to explain consciousness. Descartes was correct in thinking about unconscious mental processes in physical terms; he erred, however, in conceiving of consciousness as nonphysical. That the brain mechanisms underlying conscious experience haven't been figured out yet doesn't mean that they will remain obscure forever. In fact, recent research has begun to make some headway in understanding the brain mechanisms of consciousness, and we'll take a look at this work later in the book.

OUR BODIES, OUR SELVES

Although the mind-body problem is the favorite topic of philosophers who work in the area called the philosophy of the mind, some of these philosophers have other things on their minds. One that is particularly relevant to us, and that is closely intertwined with the mind-body problem, is the issue of what constitutes a person. Is a person a body, a mind, a mind *in* a body? Does a person have to be human? Are all humans persons? Could a creature from another planet be a person? Can a human lose personhood as a result of brain damage, insanity, or moral transgressions? When during life does personhood start and stop? Is an embryo or an infant a person? What about someone who lingers for months in a coma from which he will, by medical prediction,

never recover? The latter questions have wide-ranging social and legal implications, but only if the former ones can be answered in some reasonable way. If we can't establish precisely what a person is, it matters little whether we are one or not. John Locke had something like this in mind when he said, hundreds of years ago, that *person* is a "forensic term, appropriating actions and their merit; and so belongs only to intelligent agents, capable of a law, and happiness and misery. . . . This personality extends itself beyond present existence to what is past, only by consciousness.

Peter Strawson is perhaps the best-known modern philosopher in this area. His much-cited paper "Persons" starts with Ludwig Wittgenstein's assertion that bodies are not in possession of the states of consciousness that come out of them: "The I occurs in philosophy through the fact that the 'world is my world.' The philosophical I is not the man, not the human body, or the human soul . . . but the metaphysical subject, the limit—not a part of the world." You'll probably be happy to know that Strawson, too, was puzzled by these words, which he called impressive but obscure.

Strawson was motivated by Wittgenstein's arcane statements to try to explain the idea of something that is both a subject of experiences (that is, is conscious) and is a part of the world (that is, is dependent on a body). He wanted to understand the relation between two questions: Why do we ascribe states of consciousness to our bodies, and why do we ascribe states of consciousness to anything at all? Descartes had raised the first when he said, "I am not lodged in my body like a pilot in a vessel," and Wittgenstein the second with his statement that "The thinking, presenting subject—there is no such thing."

According to Strawson, because we can attribute our own states of consciousness to ourselves, others like us must also have similar states of consciousness. If we can figure out how to identify those who are like us, we can know to whom consciousness should be attributed—in other words, we can know who is a person. To do this, he distinguished between two kinds of statements: those that obviously can be applied to material bodies that also exhibit consciousness ("is in pain," "is thinking," "believes in God") and those that can be applied equally to material bodies that are conscious and that are not ("is heavy" "is tall," "is hard").

Like Locke and Strawson, many philosophers have taken the view that personhood is a characteristic of intelligent, conscious creatures, that consciousness is, in fact, the quality that defines personhood. But others demand more, in the form of a moral element. This was implicit in Locke, as well as in the writings of Kant. Daniel Dennett

combined the thought of Locke and Kant, and other philosophers, proposing that there are two interrelated notions of a person, one moral and one metaphysical. The metaphysical person is a thinking, feeling, intelligent, conscious agent, while the moral person is one who is accountable for his actions. Dennett asks whether being a person in the metaphysical sense automatically makes one a person in a moral sense, or does it merely make a moral capacity possible. He goes on to list several conditions of personhood. A being is a person if it is rational, verbal, conscious, and, in fact, self-conscious, capable of being acted toward in a certain way, and capable of reciprocating when acted toward in this way. Dennett's list also borrowed from John Rawls, who argued, "To recognize another as a person one must respond to him and act towards him in certain ways," and from Thomas Nagel, who affirmed that "extremely hostile behavior towards another is compatible with treating him as a person." But Dennett, in the end, concludes that these are necessary but not sufficient conditions for defining a person—that there is, fundamentally, no way to set a passing grade for personhood that is not arbitrary.

The concept of the self, which is of utmost importance to us here, is closely related to the philosophical notion of a person. Within philosophy, there has, in fact, been a growing interest in the self, an outcome of which has been the emergence of distinctions between different aspects of the self. One much-discussed distinction is between the minimum and the narrative self.' The former is an immediate consciousness of one's self, and the latter a coherent self-consciousness that extends with past and future stories that we tell about ourselves. The narrative self bears some relation to the postmodern notion that the self is socially constructed. While social construction is often viewed as diametrically opposed to a scientific view of man, the two are not necessarily at odds with each other since brains, in the end, are responsible for both the behaviors that collectively constitute the social milieu, and for the reception by each individual of the information conveyed by this milieu.

In focusing on consciousness as the leading metaphysical feature of who we are, philosophers interested in the question of personhood and the self leave out much of who we are—all the nonconscious aspects. And in dividing the world into material objects and conscious selves or persons, as Strawson did, nonhuman animals are placed in a kind of ontological limbo, since a nonhuman animal cannot be a person.

Although other animals are not conscious in the human sense, they are not simply objects, like rocks or chairs. They are living creatures with nervous systems that make it possible for their bodies to interact with and change the material world in ways that rocks and chairs cannot. The concept of a person, a conscious self, while useful as a way of evaluating issues related to being human, is thus less valuable as a general-purpose concept for understanding existence in the context of our animal ancestry. And because we must pursue many aspects of how the brain works through studies of nonhuman organisms, we need a conception of who we are that recognizes the evolutionary roots of the human body, including the brain.

Though not as widely discussed as conscious aspects of the self, nonconscious aspects are nevertheless important. They are essential to the Buddhist attempt to eliminate the conscious self, to ideas about multiple selves by William James and others as well as to notions of a primitive, nonconceptual, or ecological self that exists outside of conscious awareness. Once we accept that the self of a human can have conscious and nonconscious aspects, it becomes easy to see how other animals can be thought of as having selves, so long as we are careful about which aspects of the self we are ascribing to each species in question.

The self, then, is a notion that can be conceived of along an evolutionary continuum. While only humans can have the unique aspects of the self made possible by the kind of brains that humans have, other animals have the kinds of selves made possible by their own brains. To the extent that many of the systems that function nonconsciously in the human brain function similarly in the brains of other animals, there is considerable overlap in the nonconscious aspects of the self between species. Obviously, the more similar the brains, the more the overlap.

The extent to which other animals have any kind of consciousness is, unfortunately, impossible to know. We can speculate, but because the human mind cannot become a cat, dog, bird, lizard, frog, or fish mind, we cannot know with certainty how such a question should be answered. Descartes's greatest contribution was perhaps his conclusion that the only thing he could know with certainty was his own mind. So long as we are talking about other

animals with brains like our own (that is, other humans), we can have some confidence that their mental states are like ours. But we cannot with any degree of certainty extrapolate from our own mental states to those of other species.

In spite of having gotten this far with some key concepts from philosophy, the fact is that philosophy will probably not give us the kind of foundation we need to pursue the relation between the self and the brain.) To state that a mind, or a person, or a self is all physical, or all mental, or partly physical and partly mental, or something else altogether (like the product of socially constructed relations between people), lays out the territory in a way that is useful for analyzing broad categories of experience within and between species, but does not tell us much about how to pursue mechanisms in the brain. If we are going to figure out how it is that our brains make us who we are, we need a way of linking a fairly detailed conception of who we are to neural functions. Psychology may be more relevant to that purpose.

MIND SCIENCE

Psychology was actually a branch of philosophy until the late nineteenth century, when Wilhelm Wundt, a German physiologist, began doing experiments to understand the way the mind works rather than just speculating about it. He and his followers, known as introspectionists, took the key steps required to convert psychology into an experimental science. Their main topic of investigation was conscious experience, which they explored by examining their own experiences, attempting to break them into essential, irreducible elements.

But early in the twentieth century, some psychologists began to argue that this was no way to conduct scientific research, since one's conscious experiences can only be known personally, and cannot be verified by others. This idea caught on and eventually spawned behaviorism, which was based on the premise that a scientifically valid psychology had to focus on observable events (behavioral responses) rather than internal states. Some of its adherents were methodological behaviorists, which meant they didn't necessarily reject the existence of consciousness, but simply believed it couldn't be studied. Radical behaviorists, in contrast, actually denied that consciousness existed. For them, mental states were illusions created by tendencies to act in one way or another. Philosophers like Gilbert Ryle adopted radical behaviorism as a resolution to the mind-body problem, eliminating the mind entirely, leaving only the physical body to be explained in physical terms. Ryle called mental states "ghosts in the machine," after the "deus ex machina" of Greek tragedy, a god that was lowered onto the stage from above to solve the problems of mortals.

Toward the middle of the century, it dawned on some scientists that the operations (computations) performed by computers were not unlike what a human does when solving a problem. This notion was embraced by some

farsighted psychologists like Jerry Bruner and George Miller, and the cognitive approach to psychology, which emphasized internal mechanisms that process information, was born. This was an attractive alternative to mindless behaviorism, and eventually the cognitive movement dethroned behaviorism and brought the mind back to psychology.

The mind that returned, though, was not exactly the one that the behaviorists had disposed of. Behaviorists had objected to the emphasis of introspectionists on mental content (the experience of the color red, for example). Cognitive scientists, however, were studying mental *processes* rather than the content of consciousness. They were more concerned with how colors are detected and discriminated than in what it is like actually to experience them.

It is now widely recognized that we can have conscious access to the outcome of cognitive processes, but we are not usually aware of the processes that were involved in generating that content. Our perceptions, memories, and thoughts generally work in happy ignorance of the processes that make them possible. For cognitive scientists, and in stark contrast to Descartes, mind and consciousness are not at all the same.

The cognitive movement had a tremendous impact on psychology, but its influence did not stop there. Information-processing concepts were also adopted by workers in linguistics, anthropology, and other social sciences, as well as mathematics and physics. And just as psychologists were conceiving of minds in terms of computer operations, computer scientists and mathematicians were pursuing the notion that computers might perform mindlike

operations, an idea that led to the field of artificial intelligence (AI). Ultimately, cognitive science emerged as an interdisciplinary approach to understanding how the mind works. It came to be called “the new science of mind.”

The fact that cognitive processes are not dependent on consciousness (actually, consciousness depends on *unconscious* cognitive processes) means that the mental vs. physical dilemma does not have to be overcome in order to study the brain mechanisms of cognition. Indeed, many of the processes studied by cognitive scientists are also topics of research pursued by so-called cognitive neuroscientists. Led by breakthroughs in understanding the psychology of cognition, cognitive neuroscientists have been very successful in relating perception, attention, memory, and thinking to underlying mechanisms in the brain.

Cognitive psychology, and its sister, cognitive neuroscience, would thus seem to be taking us ever closer toward psychological and neurobiological understandings of the self. However, this is not exactly the case. Though we understand how specific cognitive processes work psychologically and neurologically, cognitive approaches fall short when it comes to explaining the self.

First of all, by its very definition, cognitive science is a science of only a part of the mind—the cognitive part—and not a science of the whole mind. Traditionally, as we’ll see in chapter 7, the mind has been viewed as a trilogy, consisting of cognition, affect (emotion), and conation (motivation). The fact that emotion and motivation are not studied by cognitive science makes sense if cognitive science is regarded as a science of cognition, but is troubling if the field is supposed to be the science of mind. A mind without feelings and strivings (the kind of mind traditionally studied in cognitive science) might be able to solve certain problems given it by a cognitive psychologist, but it doesn’t stack up well as the mental foundation of a self. The kind of mind modeled by cognitive science can, for example, play chess very well, and can even be programmed to cheat. But it is not plagued with guilt when it cheats, or distracted by love, anger, or fear. Neither is it self-motivated by a competitive streak, or by envy or compassion. If we are to understand how the mind, through the brain, makes us who we are, we need to consider the *whole* mind, not just the parts that subserve thinking.

A second shortcoming of cognitive science is that it has not grappled successfully with how various cognitive processes interact to form the mind. Considerable progress has been made in understanding how perception, memory, and thinking work, but not about how they work together. And in light of the tripartite nature of the mind, an understanding of the self is going to require that we not only figure out how various cognitive processes interact, but also that we include emotions and motivations in the mix and figure out how they interact with one another, as well as how they interact with cognitive processes. Our hopes, fears, and desires influence how we think, perceive, and remember. A science of mind needs to account for and understand these complex processes.

And, finally, cognitive science deals with the way the mind typically works in most of us, rather than the way it works uniquely in any one of us. While we all have basically the same mental processes mediated by the same brain mechanisms, the way these processes and mechanisms operate is determined by our particular genetic background and life experiences.

It would be hard to overstate the importance of cognitive science. It has been extremely successful as a research program, and has revolutionized the way we conceive of the mind. So when I single out the shortcomings of the

field, I do so not to condemn it, but instead to simply point out that it’s incomplete when it comes to understanding what makes us who we are.

THE PERSONALITY CONTEST

Psychology, as we know it today, is an imperfect marriage between two distinct approaches to the workings of the mind that emerged in the late nineteenth century. One is the experimental approach, which emphasizes the way specific mental processes, like perception or memory, typically work. This is the approach that gave rise to cognitive psychology. The other approach is more concerned with how well-adjusted people are and how they might change their behavior to improve their psychological well-being. It focuses on individuals and their idiosyncratic traits,

habits, feelings, and thoughts, rather than on the way things work in most people most of the time. The various forms of psychotherapy in use today are outgrowths of this approach, which has also been a fountainhead for theories of personality. This is the kind of psychology that is portrayed in films and novels and is what people usually have in mind when they think of what a psychologist is.

Ideas about personality are probably as ancient as ideas. Around 400 B.C., Hippocrates, for example, proposed that one's health and character were determined by the interaction among four bodily humors (blood, phlegm, black bile, and yellow bile), each of which, in turn, reflected four cosmic elements (earth, water, air, and fire). Six hundred years later, Galen expanded the theory, proposing that excess in one or another humor gave people distinctive personalities. (Excess blood led to a sanguine, enthusiastic personality; too much black bile made one melancholic; abundance of yellow bile led to irritable or choleric temperament; and overproduction of phlegm gave rise to a slow, apathetic, or phlegmatic person.)

Although views of personality, temperament, character, and the self continued to be developed over the centuries, modern approaches essentially began with Freud's psychoanalytic theory. Subsequent theories have for the most part been variations on or reactions to Freud, and fall into several broad categories. These include neo-Freudian psychodynamic theories, organismic or self theories, trait theories, behavioral or learning theories, and cognitive theories.

Personality theorists clearly have had valuable insights into the workings of the human mind, and have guided therapists in their efforts to help people adjust to life's challenges. But the various theories are often directly contradictory. There were feuds within psychoanalysis even in Freud's day (Jung, for example, broke away from strict Freudianism). Later, the neo-Freudians had disputes as well. For example, some maintained Freud's emphasis on repressed sexual urges as the root of anxiety; while others replaced sexuality with social and/or cultural factors as the core psychoanalytic concept. Psychoanalytic theory today is probably best viewed as a family of theories rather than a single well-defined view of how the mind works and how it breaks down in psychopathology.

But differences within psychoanalytic theory pale compared to differences between psychoanalytic and other personality theories. Some theories focus on psychopathology, while others are more concerned with the nature of the well-adjusted person. Unconscious motivation plays a key role in some theories, while others go in the opposite direction and focus almost exclusively on conscious strivings. Behavior is motivated in multiple ways in many theories, whereas in others, a single motive is emphasized (e.g., sexual gratification or self-actualization). Social considerations are important in some but are less crucial in others. Biological factors, especially genes, are believed to underlie stable personality traits over one's life span for some theorists, but others emphasize the role of learning and situational (especially social) factors in determining behavioral and mental states.

One possible explanation for the diversity of personality theories is that the topic is simply so difficult that no one has quite figured it out yet. Alternatively, there may have not been a clear winner in this personality contest not because the various theories proposed to date are all wrong, but because many are at least partly correct. If this is true—and I believe this is the case—then the best way to construct a view of the self might be not to pit the various theories against one another but rather to synthesize across them.

A VIEW OF THE SELF

So far, I've used the terms *personality* and *the self* rather loosely. Now it's time to get more specific. From here on, when I use the term *the self* I am referring to the totality of the living organism. This notion subsumes the idea of personality and is similar to what William James had in mind when he described the self as the sum total of who one is (see the opening paragraph of this chapter). But in order for this view of the self to be useful to our pursuit of how the brain makes the self possible, we need to refine it considerably.

In modern personality theory, as in philosophy, the notion of the self typ

ically refers to the conscious self, in the sense of having self-knowledge, a self-concept, and self-esteem; of being self-aware, self-critical; of feeling self-important; and of striving toward self-actualization. Carl Rogers, a pioneer

psychologist of the self, summed up this view, defining *the self* as “the organized, consistent conceptual gestalt composed of perceptions of the characteristics of the ‘I’ or ‘me.’” (O For Rogers, these perceptions are available to awareness, though not necessarily in awareness.” Modern self psychologists like Hazel Markus have a similar focus on self consciousness.” These psychologists do not deny that some aspects of mental life occur unconsciously, but instead minimize the importance of the unconscious components of the mind in favor of the notion of a self as an active agent in the control of mental states and behavior.

In spite of this long tradition of emphasis on the self as a conscious entity, the self that we are aware of, or can be aware of, is not the entirety of what the term *the self* refers to. The psychologist Ruth Munroe, for example, argues for a more fundamental view. She points out that “a sense of self which *develops* in the course of living is too far confused with the truly necessary organismic self.” Munroe is questioning whether the “sense of self” that develops over time is the whole self. In other words, she is arguing that the self that we are aware of and strive to improve, the self that we have a sense of, the self that many personality theorists have been enthralled with, is too narrow a view of what the self really is.

The existence of a self is a fundamental concomitant of being an animal. All animals, in other words, have a self, regardless of whether they have the capacity for self-awareness. As a result, the self consists of more than what self-aware organisms are consciously aware of. Indeed, recent research in social psychology has emphasized that many important aspects of human social behavior, including decision-making as well as the way we react to members of racial and ethnic groups, are mediated unconsciously.” These differences within organisms (conscious vs. unconscious aspects) and between organisms (creatures with and without consciousness) are not captured by an undifferentiated notion of the self, but can be accounted for by distinguishing between explicit and implicit aspects of the self.

Things we consciously know about who we are make up the *explicit* aspects of the self. These are what we refer to by the term *self-aware* and constitute what we call our self-concept; they are what the self psychologists are interested in. The *implicit* aspects of the self, by contrast, are all other aspects of who we are that are not immediately available to consciousness, either because they are by their nature inaccessible, or because they are accessible but not being accessed at the moment. All animals have implicit selves, but only animals that have the capacity for conscious self-awareness have explicit selves (this is why the existence of a personality in a pet does not necessarily mean that the pet is conscious in the human sense).

This view of the self contrasts with the idea of a person as elaborated by philosophers like Strawson, Dennett, and others. Only humans can be persons, but all animals can have selves, especially when we allow for the distinction between implicit and explicit selves. One might want to broaden the notion of a person to account for explicit and implicit aspects. This would solve one problem (the fact that there's more to a person than what that person is conscious of) but would leave another unaddressed (the relation between persons and other animals).

That explicit and implicit aspects of the self exist is not a particularly novel idea. It is closely related to Freud's partition of the mind into conscious, preconscious (accessible but not currently accessed), and unconscious (inaccessible) levels. However, Freud's terms carry much theoretical luggage that I want to leave behind.

The terms *implicit* and *explicit* are themselves not completely neutral. They are borrowed from the study of memory, where it is now widely recognized that the brain system involved in forming explicit, consciously accessible memories is distinct from a variety of other systems that are capable of learning and storing information implicitly, which is to say without conscious awareness. Actually, since most brain systems are plastic, and work outside of consciousness, they can be thought of as implicit memory systems or, better yet, as systems that are able to store specific kinds of information implicitly. To the extent that our life's experiences contribute to who we are, implicit and explicit memory storage constitute key mechanisms through which the self is formed and maintained. Those aspects of the self that are learned and stored in explicit systems constitute the explicit aspects of the self. To be self-aware is to retrieve from long-term memory our understanding of who we are and place it in the forefront of thought. In contrast, those aspects of the self that are learned and stored in implicit systems make up the implicit aspects of the

self. We use this information about our selves all the time, even though we may not be consciously aware of it. The way we characteristically walk and talk and even the way we think and feel all reflect the workings of systems that function on the basis of past experience, but their operation takes place outside of awareness. I will have much to say in later chapters about the workings of explicit and implicit memory functions of the brain.

The self is a unit, in the sense that organisms go to great pains to keep themselves alive and well. Physical damage to one's appearance is not taken lightly (remember the harm that face removal from the Egyptian tomb was meant to achieve), nor are insults to one's character. Both implicit and explicit systems are utilized to accomplish this unity in humans. But self-preservation is a universal motive, independent of whether the organism is aware that it is working toward this goal. A cockroach can scamper away when a human foot approaches without being explicitly aware of being in danger, the same way that a single-cell bacterium can detect and move away from harmful molecules in its chemical world.

The self is not static. It is added to and subtracted from by genetic maturation, learning, forgetting, stress, aging, and disease. This is true of both implicit and explicit aspects of the self, which may be influenced similarly or differently at any one point. For example, a mild compliment may only be registered and stored in explicit memory, but glowing praise, registered explicitly, might lead to the arousal of emotion systems that then also store aspects of the experience implicitly. On the other hand, stress is known to impair explicit memory while at the same time enhancing the implicit memory functions of emotion systems.

As important as learning is, not all aspects of the self are learned. Some are due to our genetic heritage. All of the capacities that we have as *Homo sapiens*, including our capacities to learn and remember, are made possible by the genetic makeup of our species. What we place in our individual memory systems is a function of our unique experience, but the existence and basic mode of operation of these systems are due to our species's genes. At the same time, we each have a family genetic history that is a variation on the theme of being a human, and a personal set of genes that is a variation on our family's, and these variations also influence who we are.

The most well-articulated view of the role of genes in shaping behavioral and mental characteristics comes from biological trait theories of personality, which propose that one's enduring qualities are due to one's genetic background. Considerable evidence has been amassed to support the view that some traits, such as the extent to which one is extroverted (gregarious) vs. introverted (shy, fearful, withdrawn), are highly influenced by one's genetic history. Nevertheless, there are two important caveats to genetic theory of personality

First, genes have been found to account for only about 50 percent of a particular personality trait. What this statement means is that genes account at most for half of a given trait, not that half of all of personality is accounted for by genes. For some traits, genetic influence is *far* less and is often not measurable. Introversion is probably the trait with the strongest genetic influence. Although many extremely shy, introverted children tend to become anxious, depressed adults, some do just fine. Is this because the genetic influence in the latter group was temporary, or because the genetic tendency was squelched? The fact that when extreme introversion is caught early, it can be reversed to some extent by a supportive family environment suggests that genes do not fully dictate psychological destiny. Life's experiences, in the form of learning and memory, shape how one's genotype is expressed. Even the most ardent proponents of genetic determination of behavior admit that genes and environment interact to shape trait expression. It's a matter of how much, not whether, both contribute.

The second caveat to the genetic account of personality stability comes from research showing that people are not always true to their so-called personality traits. One may be shy at work or in social groups, for example, but domineering at home. In fact, when psychologists have examined the consistency of behavior across situations, the results have not supported the view that people act consistently in different situations. Observations such as these suggest to Walter Mischel that behavioral and mental states are not dictated by constitutional factors but instead are situationally determined. Mischel argues that the ability to predict behavior depends upon knowing about a person's thoughts, motivations, and emotions relative to a particular set of circumstances. He describes these as "if... then relations." "If" you are in situation A, "then" you do X, but "if" in situation B, "then" you do Y. According to Mischel,

people don't possess stable personality traits over time, but stable "if. . . then" profiles.

As with most polarized arguments in psychology, there is truth in both the situational and the trait views. The stronger the genetic contribution to a particular characteristic, the more likely it will be expressed uniformly in different situations. At the same time, situations vary in the extent to which they dictate the way we act. A red traffic light will cause most people to stop, regardless of whether they are generally aggressive or timid, whereas a yellow light allows more latitude for tendencies like aggression or timidity to be expressed.' We'll visit questions about genetics and personality again in chapter 4, when we explore how the brain is built.

In proposing that the self exists, I run the risk of reifying something that is, ultimately, not real. Bob Dylan, for example, said, "I change during the course of a day. I wake and I'm one person, and when I go to sleep I know for certain that I'm someone else. I don't know who I am most of the time. It doesn't even matter to me." And, according to Philip Roth, "All I can tell you with certainty is that I, for one, have no self, and that I am unwilling or unable to perpetrate upon myself the joke of a self." Mark Epstein, who has tried to integrate psychoanalysis and Buddhism, points out that the ego's image of itself (its object image) is always lacking as an account of the subject (the self), implying that much of the self is, in essence, implicit. While the whole self is not usually encountered by the individual who possesses it (who *is* it), or by others, it nevertheless exists.

'What then is it? In my view, the self is the totality of what an organism is physically, biologically, psychologically, socially, and culturally. Though it is a unit, it is not unitary. It includes things that we know and things that we do not know, things that others know about us that we do not realize. It includes features that we express and hide, and some that we simply don't call upon. It includes what we would like to be as well as what we hope we never become.

The fact that all aspects of the self are not usually manifest simultaneously, and that different aspects can even be contradictory, may seem to present a hopelessly complex problem. However, this simply means that different components of the self reflect the operation of different brain systems, which can be but are not always in sync. While explicit memory is mediated by a single system, there are a variety of different brain systems that store information implicitly, allowing for many aspects of the self to coexist. As William James said, "Neither threats nor pleadings can move a man unless they touch some one of his potential or actual selves." In *Orlando*, Virginia Woolf pointed out, "A biography is considered complete if it merely accounts for six or seven selves, whereas a person may well have as many thousand." Or as the painter Paul Klee expressed it, the self is a "dramatic ensemble."

THE SELF AND THE BRAIN

Theories of the self and personality are not usually framed in ways that are compatible with our understanding of brain function. How, then, can we relate the complex constellation I've called the self to the systems and synapses of the brain?Bo The goal of the rest of the book is to answer this question. However, a brief preview is in order.

The self can be understood in terms of brain systems involved in learning and storing information, in explicit and implicit systems, about things that are significant in people's lives. The processing by these systems always occurs in a physical and social context (a situation) and is performed by networks that function the way they do because of both genetic inheritance and past experiences. Put this way, in order to understand the self, we need to explain how brain systems underlying thinking, emotion, and motivation (the men— tal trilogy) develop under the influence of nature and nurture, and how these systems make it possible for us to attend to, perceive, learn about, and store and retrieve experiences. We especially need to explain how different systems interact with and influence one another. Without these interactions, and the mental integration they engender, each of us would simply be a collection of isolated mental functions rather than a coherent person.

The point, though, is not simply to state that learned and innate interactions between cognitive, emotional, and motivational processes make us who we are, but instead to explain *how* these interactions work. And the explanation that I will pursue in the remainder of this book involves neural, especially synaptic, mechanisms. I believe,

in short, that an answer to the question of how our brains make us who we are can be found in synaptic processes that allow cooperative interactions to take place between the various brain systems that are involved in particular states and experiences, and for these interactions to be linked over time. It is probably not at all obvious what this statement means at this point in the book. Before it will begin to make sense, we need to cover more ground.

CHAPTER THREE

THE MOST UNACCOUNTABLE OF MACHINERY

MY OWN BRAIN IS TO ME THE MOST UNACCOUNTABLE OF MACHINERY— ALWAYS BUZZING, HUMMING, SOARING ROARING DIVING, AND THEN BURIED IN MUD. AND WHY? WHAT'S THIS PASSION FOR?

—Virginia Woolf

Most of us are as mystified by our brains as Virginia Woolf, though perhaps less eloquent in our ignorance. Still, everyone has heard a few things about the wrinkled blob in the noggin—for instance, that we use only 10 percent of it. But who came up with this number? And why would we even have the rest if it weren't useful? Evolution doesn't usually make organs in such a way that they mostly go unused, just in case someone figures out one day what to do with the extra material. It's hard to imagine how 90 percent of the brain, lacking in value for most of us most of the time, could have ever come into existence. Researchers have been looking into what the brain does for many years now, and from what they have discovered, it doesn't seem that most of it is, in fact, resting idly.

People also tend to carry around with them one or both of two additional erroneous beliefs about the brain. The first is that functions of the brain, like perception, memory, or emotion, are located in specific areas. The other is that chemicals floating around in the brain determine our mental states. Unlike the 10 percent myth, these are actually part truths that, taken out of context, are patently false. We know, at least in a general sense, how the brain works, and it's not by islands of brain tissue or by isolated chemicals operating independently. Particular areas are important, but not on their own: they participate in functions by way of their synaptic connections with other areas. Chemicals *are* also important, but mainly because of their work at synapses within functional systems.

This chapter will give an account (albeit an abbreviated one) of this “most unaccountable of machinery,”

describing some basic facts that are necessary to understand the brain's synaptic systems. Although the discussion will have to get a bit technical along the way, this information is essential to my attempt to relate the self to synapses. Because I've kept things simple, those already in the know may wish to skip ahead. However, the novice will get a crash course on what neurons are, how synapses connect them together, and why synaptic connections are the key to brain function.

BRAINS: SO DIFFERENT, YET ALL THE SAME

We mammals belong to the group of animals called vertebrates, a subphylum we share with other backboned creatures, including birds, reptiles, amphibians, and fish. Mammals and birds separately descended from reptiles millions of years ago. In spite of this common ancestry, the brains of reptiles, birds, and mammals look very different. Beneath these dissimilarities, though, there's a common plan that's rigorously adhered to.

Every vertebrate brain can be divided into three broad zones: the hind-brain, midbrain, and forebrain. In the early years of the twentieth century, neuroscientists discovered that damage to each zone had a different predictable consequence.' For example, in studies of cats, it was found that purposeful, voluntary behavior and problem-solving ability were impaired when the forebrain was damaged. Nevertheless, even with massive injuries to the forebrain, some semblance of normal coordinated behavior remained. Such compromised animals could orient toward a noise or withdraw their paw from heat, and could walk, eat, and groom. They could even display full-blown emotional responses, especially those typically expressed in anger or fear, if the hypothalamus, a small region situated at the base of the forebrain, was spared. When larger lesions were made that removed all of the forebrain, including the hypothalamus, only rudimentary responses remained. These animals, when challenged with intense stimulation, could hiss, bare their teeth, unsheathe their claws, or swipe a paw, but could not manage to put all of these behaviors together into a coordinated defense or attack response. When the midbrain was damaged, the animal was essentially comatose—

alive physically, but not behaviorally or psychologically. And when the hind-brain was destroyed, life itself ceased.

From these crude experiments, it was concluded that the hindbrain controls very basic functions, those necessary for staying alive; the midbrain is involved in maintaining wakefulness and coarse, isolated behavioral reactions; and the forebrain coordinates complex behavioral and mental processes. It should not be surprising, given these effects of brain damage, that the forebrain (necessary for thinking and problem-solving) is the region that differs the most between mammals and other vertebrates and the hindbrain (necessary for life) the least. Nevertheless, all three levels are represented in all vertebrates, and even the evolutionarily advanced forebrain is structured according to a common underlying organizational plan that is applicable to every vertebrate species.

For example, the human forebrain consists of several subdivisions, one of which is the wrinkled outer layer, the neocortex. This is the part of the forebrain that makes possible many of our higher mental functions. The designation *neo* reflects the fact that this brain region was, for many years, believed to be evolutionarily new, having emerged when mammals evolved from reptiles. Other vertebrates were thought to have a primordial or older cortex but not a mammalian or neocortex. This view began to change, though, in the late 1960s and early 1970s, when new techniques for studying the brain became available. Based on the patterns of chemical staining and nerve connections discovered with these techniques, the organization of the brain came to be better appreciated, and researchers were able to use this information to find the equivalent (or at least the semblance) of a neocortex in both birds and reptiles, suggesting both that it wasn't so new after all and that it certainly wasn't unique to mammals. The reason this cortex had not been found in these animals earlier was because of its unusual location, buried beneath other brain areas, instead of resting on top, as it does in mammals.

While at the level of overall brain structure a similar organizational plan applies to many different animals, it is not the case that all brains are the same. A given brain area can vary enormously in size and complexity between different species, allowing some animals to do things that others cannot. In amphibians, for example, an area in the

midbrain called the rectum is especially well developed, making it possible for most frogs to thrust their tongue into the flight path of an insect and capture it, a feat most people can't accomplish. Bats and rats can hear things that we cannot, and bees use a magnetic sense, which we do not have, to guide their movements. Different species have been subjected to different evolutionary pressures and their brains reflect their unique histories.⁷

The most obvious difference between the mammalian and other vertebrate brains is the extent to which the cortex has expanded. Although, as we have seen, reptiles and birds are now known to have some neocortex, the mammalian neocortex is far more elaborate than the equivalent areas in these other species. And within mammals, there are distinctions as well: the neocortex is bigger and more differentiated in primates than in rodents, and in humans more so than in monkeys. These changes in cortical size and complexity are, however, superimposed on a basic neocortical plan. For example, in all mammals, processes related to sensation (vision, audition, touch) are represented in the rear and processes involved in controlling movement in the front of the cortex.

Within a given species, the similarities of cortical organization are striking. Early anatomists discovered that the major patterns of cortical wrinkles, which appear to be randomly arranged to the uninitiated eye, are amazingly consistent from person to person, and can be used as landmarks to identify various regions of the neocortex. What's remarkable is that these purely structural parcels, defined by the wrinkles, turn out to correspond to functional divisions, areas that participate (by way of their synaptic connections with other cortical and/or subcortical areas) in different aspects of mental life and behavior.” For example, the area of the cortex involved in controlling precise movements of various body parts is located just in front of the central sulcus, one of the major wrinkles in the cortex, while touch, hearing, and visual areas are defined by their own wrinkles, as are areas involved in language comprehension and speaking. On careful examination, some variation in the organization of cortical or other brain areas is evident in different people, but the basic overall architectural plan of the brain is pretty much the same in any two individuals.

In spite of the tremendous similarity of our brains, we all act differently, have unique abilities, and have distinct preferences, desires, hopes, dreams, and fears. The key to individuality, therefore, is not to be found in the overall organization of the brain, but rather in the fine-tuning of the underlying networks. To understand the defining qualities of each person, we need to go beyond the superficial organization of the brain (its division into broad regions and areas within these) and turn to the microscopic structure and function of neural systems, and especially to the cells and synapses that constitute them.

THE CELL WAR

All organs and tissues of the body are composed of cells. But unlike the cells in other body parts, brain cells, or neurons, directly communicate with one another. There's nothing magical about the process—neurons are simply built in a way that allows them to exchange information with one another in ways that other cells cannot.” Common patterns of communication between neurons ensure that all human brains work in basically the same way, whereas subtle differences in these patterns of communication give rise to the distinctive qualities that we each have.

The existence of cells in the brain and other parts of the body is taken for granted today, but this knowledge was only made possible by the further development of the microscope in the nineteenth century. Around 1837, Matthias Schleiden, a German botanist, first proposed that plants were made up of discrete units, or cells. The following year, his friend Theodor Schwann extended the notion to animals, and thereby brought botany and zoology together in a single theory, the so-called cell theory,⁸ which argued that all living things are composed of cells.

Whether cell theory was applicable to the brain was a topic that was fiercely debated for decades. When early brain anatomists examined brain tissue under a microscope, they did see structures resembling cells. But unlike cells in other organs, brain cells had fine fibers extending out of their. Some scientists concluded that this meant that the brain

unique—not composed of discrete cells but instead made up of an entangled mesh or reticulum of continuously connected elements. Others, though, argued that the fundamentals of cell theory applied equally to the brain.

Two of the major figures in the debate were Santiago Ramón y Cajal of Spain and the Italian anatomist Camillo Golgi. Golgi, working in his kitchen, invented methods for staining the brain that allowed better visualization of its microscopic anatomy. He favored the reticular theory. Ironically on the basis of the methods pioneered by Golgi, Cajal argued forcefully for the application of the cell theory to the brain, and won many converts. One of these was Wilhelm Waldeyer, who in 1891 published a paper in which he suggested that brain cells be called neurons. In this paper, he also coined the phrase *the neuron doctrine* to account for the application of the cell theory to the brain. Cajal apparently considered the doctrine his, at least in spirit if not name, and was not happy to have had his thunder stolen by Waldeyer. But the loss in stature, if any, was temporary. Every graduate student in neuroscience today knows who Cajal was, whereas few have ever heard of Waldeyer.

One of the early and largely unrecognized soldiers in the neuron war was the young Sigmund Freud. After completing his medical training in Vienna, Freud accepted a position as a *famulus*, or research scholar, and studied the nervous system of fish and crayfish. As early as 1883, long before the neuron doctrine was codified, he promoted the idea that nerve cells are physically separated from one another. This concept later figured prominently in one of his earliest forays into psychological theory. In *Project for a Scientific Psychology*, written in 1895 but unpublished for many decades, Freud stated that “the nervous system consists of distinct and similarly constructed neurones. . . which terminate upon one another.” He introduced the term *contact barriers* to describe the points where neurons abut, and suggested that interactions between neurons across contact barriers make possible memory, consciousness, and other facets of the mind. Although these notions were amazingly sophisticated for their time, Freud felt that progress in understanding the brain would be too slow for his taste and so abandoned a neural theory of the mind in favor of a purely psychological one. The rest is history.

Two years after Freud wrote his *Project*, Sir Charles Sherrington proposed a different term for the connections between neurons. Sherrington had been working on the reflex problem. A reflex is the simplest kind of neural circuit that controls behavior. When your physician taps you on the knee, your leg jerks because the tap elicits sensations that are transmitted along *sensory* nerves that originate in your knee and travel to your spinal cord. The messages in the sensory nerves trigger activity in *motor* nerves that come out of the spinal cord and end in your leg muscles, leading to the jerk. Sherrington realized that the gap between the sensory and motor neurons had to be bridged somehow if information carried by the sensory nerves was to be transferred to the motor nerves. He was probably unaware of Freud’s contact barriers, and chose to call the gaps *synapses*, derived from the Greek word meaning to clasp, connect, or join. The notion of synapses as points of communication between cells is one to which we still adhere, and which is essential to our efforts to understand who we are in terms of brain mechanisms.

In 1906, Cajal and Golgi shared the Nobel Prize for their groundbreaking research on brain anatomy. Although the neuron doctrine had gained considerable support by then, Golgi clung bitterly to the reticular theory at the award ceremony. Still, definitive proof that the nervous system is composed of cells did not come until many years later. With the invention of the electron microscope in the 1950s, scientists could finally examine the brain in sufficient resolution to see that the tiny fibers extending out of a neuron do not typically make direct physical contact with neighboring cells. Indeed, they are separated by tiny spaces, synaptic spaces, across which the brain does its business.

WHAT MAKES NEURONS SPECIAL?

By knowing the function performed by a few cells of most organs in the body, whether the liver, kidney, or gall bladder, you can deduce the organ’s overall function. This is not true of the brain, however, where cells participate in myriad activities, from seeing and hearing to thinking and feeling, from awareness of self to the incomprehension of infinity. The architecture of a neuron helps us begin to understand why the brain is so multifunctional, while organs like

the pancreas and spleen are not.

Neurons have two major parts. The first is the cell body, which is involved in important housekeeping functions, such as storing genetic material and making proteins and other molecules that are necessary for the cell's survival. The cell body does much the same work in neurons as it does in other cells. The major structural difference between neurons and other cells lies in the special appendages that neurons have—the nerves. These fibers, which extend out of the cell body, are what caused all the confusion in the nineteenth century about whether the brain was, like other organs, composed of discrete cells.

Nerve fibers are sort of like telephone wires. They allow neurons in one part of the brain to communicate with neurons in another. By way of these connections, communities of cells that work together to achieve a particular goal can be formed across space and time in the brain. This capacity underlies all of the brain's activities and is absent in other organs.

There are two varieties of nerve fibers, axons and dendrite. Axons are output channels, and dendrites are input channels. An axon carries messages to other cells. It can end nearby, allowing communication with its close neuronal neighbors, or it can stretch over very long distances, as much as several feet. If you are standing still and decide to take a step, the movement of your leg on the basis of your decision involves axons that originate in cell bodies located in the movement control regions in the frontal cortex (just behind your forehead) and that travel uninterrupted to the base of the spinal column (in the region of your lower back).

The end of the axon, called the terminal, is the point at which the sending neuron communicates with receiving neurons. Although terminals most often form connections with dendrites, they can also contact cell bodies or other axons. Dendrites, too, sometimes communicate between one another. In order for the long axons descending from your frontal cortex to your spinal cord to cause your leg to move, the terminal has to contact dendrites of the receiving cells in the spinal cord. The axons of these receiving cells then extend out and terminate at muscles in your leg. The arrival of signals at the muscle leads to contraction, and thus movement.⁷

Many dendrites have little knobs called spines extending from them. These are readily seen when brain tissue is stained with the methods developed by Golgi. Spines are especially important as receivers of messages from axons, and play a key role in brain development, as well as in learning and memory; as we will see later.

Most neurons have only one axon. However, each axon branches many times before it ends, allowing a single neuron to spawn many terminals. The result is that the messages sent out from one cell can affect many others. This is called divergence. At the same time, each neuron can receive inputs from numerous others. This is called convergence.

The point at which the sending and receiving elements of neurons meet is our star, the synapse. Because information usually flows across the synapse starting from the axon terminal, this side is said to be presynaptic, and the receiving side, often occupied by a dendritic spine, postsynaptic. As Sherrington noted, because a synapse is a space between the sending and receiving cells, something has to cross the synaptic space in order for the two cells to communicate.

GALVANIZED FROG LEGS

The question of how information is exchanged between neurons across the synaptic space is tied in closely with that of how information is transferred along a nerve fiber of a single neuron. Before we consider synaptic transmission, we therefore need to consider nerve conduction.

In the 1770s, Anton Mesmer, a Viennese physician, had been using iron magnets to treat a variety of physical and mental maladies, until he found he could have the same effect without the magnets when he looked into a patient's eyes and waved his hands over the afflicted body part. This was the birth of mesmerization, or hypnosis. Mesmer believed that some mysterious, magnetically sensitive fluid was present throughout the universe, including the human body, and that he could help his patients by using his own animal magnetism to alter this fluid. At the

time, little was known about the physiology of the nervous system, and any theory, including one as wacky as animal magnetism, seemed possible.

A few years later, the Italian Luigi Galvani noticed that the amputated leg of a frog hung from an iron trellis with a brass hook twitched during a lightning storm. He also found he could make a frog leg kick at any time he wished if he touched the nerves within the wound with one metal and the foot with another. This was, in effect, the first battery. Galvani, in the tradition of Mesmer, concluded that the metals were conducting vital spirits from the frog. So-called animal electricity was an occult rather than a scientific phenomenon.

Several decades later, Carlo Matteucci, another Italian, made the first measurements of genuine electrical activity in nerves. In Germany, Johannes Muller and Emil Du Bois-Reymond, realizing the importance of this observation, began a research program that rescued electrical conduction in nerves from the world of mysticism and turned it into a thriving scientific research field.

At the time, the assumption was that nerves conducted electricity like wires. But one of Du Bois-Reymond's students, Hermann von Helmholtz, did an experiment that suggested otherwise. He calculated the speed of electrical conduction in frog nerve fibers by measuring how much time elapsed before a given muscle twitched when nerves of different lengths were electrically stimulated. Although conduction time was fast—about 40 meters per second (roughly 40 mph)—it wasn't as fast as electricity, which can under certain conditions flow through a wire at about the speed of light.

From these simple but informative experiments it became clear that nerves do conduct electricity, but in a special way. Electricity does not flow passively through a nerve as it does through a wire. Rather, impulses conducted through nerves are *biologically* propagated, moved along by electrochemical reactions, a process that takes a lot longer than passive physical conduction.

The biologically propagated impulse in a nerve is called an action potential. This dramatic electrical event is normally initiated at the point where the axon emerges from the cell body. Once triggered, it travels like a rolling wave down the axon toward the terminal. The propagation occurs as a kind of neurodomino effect—an electrical change in one part of the axon membrane produces a similar change in adjacent parts, and so on, all the way down to the terminal. Action potentials can be triggered artificially by electrical stimulation, which makes them easy to study, but normally they occur in a cell when orders come from synaptic inputs.

Work by many pioneering neuroscientists established the basic principles of electrical propagation in axons, which became the foundation for much of what we now know about the working of neurons. A good deal of this research was performed using the giant axons of squids, the sheer size of which made it easier to investigate electrical conduction. Especially noteworthy were the studies performed in the 1940s by Alan Hodgkin and Andrew Hux-

ley in England. Building on Ohm's law of electricity (which states that voltage is equal to current times resistance), they characterized in precise mathematical form the basic features of electrical transmission in axons. The Hodgkin-Huxley equations are still used today to calculate current, voltage, and resistance in axons.

SYNAPTIC CHATTER

The existence of electrical conduction in nerves suggested to late-nineteenth-century scientists that electrical impulses played a critical role in the normal functions performed by the brain. A key related question was whether electrical propagation was sufficient to explain how the brain worked. Sherrington's studies of reflexes determined that it was not.

Electrical impulses in sensory and motor nerves clearly seemed involved in reflexes: when sensory nerves detect a tap on your knee, they conduct electrical impulses that, in turn, lead to electrical impulses in motor nerves, and to the jerk. But how does the sensory nerve communicate with the motor neuron? Sherrington demonstrated that while electrical stimulation of a sensory nerve elicited an electrical response in the motor nerve, stimulation in the motor nerve did not evoke a response in the sensory nerve. He concluded that the junction between cells, the synapse, had a valvelike property—it only transmitted in one direction, from sensory to motor nerves. This was

particularly significant ammunition against the reticular theory, for if neurons were continuously connected and communicated only by electrical conduction, then motor nerves should have as sizable an effect on sensory nerves as the other way around. Neurons must therefore communicate with one another by some means other than mere electrical conduction.

Subsequent research revealed that the one-way conduction between neurons is due to the fact that synaptic transmission involves the release of chemicals from storage sites in the presynaptic axon terminal. These molecules are released when action potentials propagated from the cell body reach the terminal. The released chemicals then drift across the liquid-filled synaptic space and come in contact with spines or other portions of the postsynaptic cell. Because the chemical storage sites usually are present in the presynaptic terminal and not in the postsynaptic dendrite, transmission only occurs in one direction. These chemicals are called neurotransmitters, since they allow neurons to communicate across the synaptic gap—they transmit between neurons.

The chemical nature of neuronal transmission was suspected from studies in the early 1900s showing that the effects of electrical stimulation of nerves could be mimicked or blocked by certain chemical agents. But it was an ingenious experiment by Otto Loewi in the 1920s that provided the ultimate proof. He removed the hearts from two frogs, leaving the nerves connected to one heart but not the other, and infused each with a saltwater solution (similar to normal body fluid). He then electrically stimulated the nerves on one, which changed the beat rate of the heart (the heart is postsynaptic to these nerves). When he removed the solution from the stimulated heart and injected it into the other, the heartbeat changed in the unstimulated heart, much as if it had been stimulated, indicating that some chemical that had been released in the stimulated heart was transferred in the solution to the other.

While Loewi's experiments involved the connection between a nerve and a muscle—the heart, in this case—essentially the same thing happens when the connection is between two neurons. That is, the arrival of the action potential in the presynaptic terminal leads to the release of neurotransmitter into the synaptic space.

The release of neurotransmitter molecules from the presynaptic terminal is a means, not an end. Its goal is to generate an electrical response in the post-synaptic cell. Although it is often the dendrites that are the postsynaptic beneficiaries of the chemical message, the electrical change produced in the dendrite has to be propagated to the cell body, and then to the axon, before an action potential can occur. This is so because the action potential is generated in the initial part of the axon where it connects with the cell body.

The arrival of transmitter from a single presynaptic terminal is typically not sufficient to produce an action potential in the postsynaptic cell. Only if the postsynaptic cell is bombarded with transmitter molecules from many presynaptic terminals at about the same time—within milliseconds—will an action potential result).

A given postsynaptic cell is believed to receive relatively few synaptic contacts from any one presynaptic neuron. As a result, much of the convergence that drives a postsynaptic cell toward action potentials comes from the convergence of different presynaptic cells onto the postsynaptic neuron (that is, the near-simultaneous arrival of neurotransmitter from different presynaptic neurons). In order for the inputs to arrive in the postsynaptic cell body at about the same time, action potentials have to have been triggered in the various presynaptic cells at about the same time. The timing has to be adjusted for different lengths of axons, since, as Helmholtz demonstrated, the longer the axon, the longer it takes for the action potential to travel down it. Keeping time in the nervous system is a very complex job.

Once the postsynaptic cell generates an action potential, its role shifts from that of a receiver to a sender. It now becomes a presynaptic neuron that helps fire action potentials in other cells.

The full sequence of communication between neurons is thus usually electricalchemical electrical *electrical* signals coming down axons get converted into *chemical* messages that help trigger *electrical* signals in the next cell. There are also synapses through which communication between presynaptic and postsynaptic sites is purely electrical, but chemical transmission is the more prevalent form. Thus, much of what the brain does involves electricalto chemical. to electrical coding of experience. As hard as it may be to imagine, electrochemical

conversations between neurons make possible all of the wondrous (and sometimes dreadful) accomplishments of human minds. Your very understanding that the brain works this way is itself an electrochemical event.

FROM CELLS TO CIRCUITS AND SYSTEMS

Every human brain has billions of neurons that together make trillions of synaptic connections among one another. Chemicals are oozing and sparks flying constantly, during wakefulness and during sleep, during thoughtfulness and during boredom. At any one moment, billions of synapses are active.

Imagine a large cocktail party at which hundreds of people are standing around and chatting with one another. If you were to place a microphone in the chandelier at the center of the room high above the crowd, you probably wouldn't be able to make out what was being said, for the many unrelated conversations would blend together in the microphone. You'd learn more by listening in on small groups than by eavesdropping on the whole room at once. In the same vein, it's not particularly instructive to ask what all of the brain's billions of neurons and trillions of connections are up to collectively at any one time. Different groupings of cells are doing different things, so attempting to take a reading of them all together doesn't tell you much. It would be more informative to examine the operation of specific circuits or systems.

A *circuit* is a group of neurons that are linked together by synaptic connections. A *system* is a complex circuit that performs some specific function, like seeing or hearing, or detecting and responding to danger. Seeing, for example, involves the detection of light by circuits in the retina, which sends signals, by way of the optic nerve, to the visual thalamus, where the visual information is processed by circuits that relay their output to the visual cortex, where additional circuits engage in further processing, ultimately creating visual perceptions. The visual system, like other brain systems, can thus be thought of as a series of hierarchically arranged circuits linked together by synaptic connections to perform some function.

Synaptic interactions between two types of neurons, called projection neurons and interneurons, are key to understanding how circuits and systems function. Projection neurons have relatively long axons that extend out of the area in which their cell bodies are located. In a hierarchical circuit, their main job is to turn on the next projection cell in the hierarchy. They do this by releasing a chemical transmitter that increases the likelihood that the postsynaptic cell, the next projection cell, will fire its own action potential. Projection cells tend to activate or excite postsynaptic cells.

Interneurons, also called local circuit cells, send their short axons to nearby circuit, often projection neurons, and are involved in information processing within a given level of a hierarchical circuit. One of their main jobs is to regulate the flow of synaptic traffic by controlling the *activity* of projection neurons. Inhibitory interneurons release a transmitter from their terminals that decreases the likelihood that the postsynaptic cell will fire an action potential. These neurons play an important role in counterbalancing the excitatory activity of projection cells.

Projection cells tend to be idle in the absence of inputs from other projection cells. Inhibitory interneurons, though, are often tonically active, which means they are firing all the time. Part of the reason why projection cells are inactive when not being stimulated is that they receive tonic inhibition from

interneurons. As a result, when excitatory inputs try to turn on a projection cell, preexisting inhibition of the projection cell has to be overcome. The balance between excitatory and inhibitory inputs to a neuron determines whether it will fire.

The amount of inhibition affecting a cell can change from moment to moment, depending on other factors. For example, when projection cells in one area of a hierarchical circuit send enough convergent inputs at about the same time to activate projection cells in the next area, the level of inhibition in the second area usually goes up as well. This happens because the excitatory inputs to an area often activate interneurons as well as projection neurons. The momentary increase in excitatory inputs to interneurons leads to a momentary increase in their inhibitory behavior, which in turn produces a momentary inhibition of the projection neurons. So-called elicited inhibition

contrasts with tonic inhibition. Because rapidly changing states of excitation and inhibition direct the flow of traffic through the brain, it's easy to understand how a breakdown in the flow of impulses could lead to neural gridlock.

Consider an example that will help illustrate how elicited and tonic inhibition regulate excitation. Imagine a circuit consisting of two projection neurons (A and B) linked together in a series. 'When A is active, B fires. If the job of the circuit were to make B fire action potentials as often as possible as long as A is active, these two neurons would be sufficient to do the job. But suppose its job instead is to take a barrage of action potentials in A and turn them into fewer action potentials in B, something that actually occurs quite often in the brain. This could be achieved by giving neuron B an inhibitory playmate (I). This local circuit neuron, like B, receives the output of A and then connects with B. So when A fires, it turns on B and I, and each produces an output. The output of B helps turn on the next cell in the circuit while the output of I turns B off. As a result, B now produces fewer action potentials when it is fired by A.

Now suppose that the interneuron I is constantly inhibiting the projection cell B. With this tonic inhibition added in, it is going to be much harder for the input from A to trigger the projection cell. If we put more excitatory neurons in with A to drive B, and time arrives just so, the tonic inhibition can I overcome. The cell can now be continuously activated. But being stuck fast-forward is not good for neurons, which can be damaged or even destroyed by unchecked excitation. Each burst of excitation thus needs to be countered with another round of inhibition. That's where elicited inhibition like that described above, comes in. When an excitatory surge overcomes tonic inhibition, elicited inhibition can rein in the excitation, resetting the circuit, preparing it for new inputs. There are many variations on this theme of tonic and elicited inhibition, but the scenario just described gives an idea of how inhibition in general works.

Inhibition is a very useful device in neural circuits. It adds tremendously to the specificity of information processing, filtering out random excitatory inputs, preventing them from triggering activity. Only if the excitatory inputs arrive simultaneously can they overcome the inhibition and elicit activity. And once activity is elicited, inhibition is important for keeping the excitation in check and resetting the circuit.

Although many local circuit cells are inhibitory, some are excitatory. Just as inhibitory interneurons can be thought of as filters, excitatory interneurons can be viewed as amplifiers. Again imagine a circuit consisting of neurons A

and B connected in series. As before, B is associated with an interneuron, but in this case it's an excitatory interneuron (E), and the axon of A branches and contacts both B and E. When A turns on B, the interneuron E is also activated, and its output causes further excitation of B. As a result, the output of B is amplified by an excitatory interneuron just as it was reduced by the inhibitory interneuron. But, as we've seen, all this excitation ultimately has to be regulated both to maintain normal functions and to prevent injury.

THE CHEMICAL BROTHERS

The job of a projection neuron, as we now know, is to turn on the next projection cell in the circuit. This means that action potentials in the axons of projection cells have to trigger the release of chemicals that cross the synapse and contribute to the firing of an action potential in the postsynaptic cell. Projection cells thus need to use a chemical neurotransmitter that has two properties. The transmitter first must be able to act quickly at postsynaptic sites—otherwise, our perceptions and other mental states could not keep up with rapidly changing events. And it must also be able to change the electrical state of the postsynaptic cell in such a way that the occurrence of an action potential is more likely to occur. Both requirements (speed and excitation) are fulfilled by the amino acid neurotransmitter glutamate, which is the main transmitter in projection neurons throughout the brain.

Glutamate actually has two roles in body function. In addition to serving as a neurotransmitter in the brain, it also plays a major part in basic life-sustaining metabolic processes that go on continuously throughout the body. For example, it is a building block in the construction of peptides and proteins, which are basic ingredients of living tissues. And, in the brain, it helps detoxify ammonia, which is a natural by-product of certain chemical reactions. Although glutamate is now known to be a ubiquitous excitatory transmitter in the brain, its role in transmission was for

a long time hard to dissociate from its so-called metabolic functions.

In contrast, inhibitory neurons, especially inhibitory interneurons, often release the amino acid GABA (short for gamma-aminobutyric acid) from the terminals of their short axons. In contrast to glutamate, this inhibitory transmitter reduces the likelihood of an action potential being generated in the postsynaptic cell. By sending axons to nearby projection neurons, GABA interneurons thereby regulate the flow of traffic through a given area.

GABA actually was identified as a neurotransmitter long before glutamate. Because it was well established that glutamate was one of the essential chemical components involved in the synthesis of GABA, its metabolic role in GABA production hampered the discovery that glutamate was itself a neurotransmitter.

Glutamate and GABA are together responsible for much of the neurotransmission business in the brain. If you understand the work done by these two chemicals, you will understand quite a lot about how synapses function. These and all other transmitters work by attaching to molecules called receptors on the postsynaptic cell. Receptors selectively recognize and bind (literally, hold on to) transmitter molecules. Glutamate receptors recognize and bind glutamate, but ignore GABA; GABA receptors are just as selective. How, then, does the binding of glutamate and GABA molecules to their receptor molecules lead to excitation and inhibition?

All cells in the body are completely enclosed by a membrane, which defines the boundary of an individual cell. The membrane is like a formfitting bag, a spandex suit, in which the cell is contained. In neurons, it covers the axons and dendrites as well as the cell body. The space outside the membrane between neurons is called the extracellular space. The fact that the extracellular space is filled with liquid has two important consequences.

First of all, this liquid is a medium that allows transmitter molecules to cross the extracellular space between presynaptic and postsynaptic sites—the synapse. Transmitters do this by diffusing out from the terminal. The distance they have to travel is very small (it's measured in tiny units called angstroms, one of which equals one ten-millionth of a millimeter), making the postsynaptic site a close and easy target.

The second point is that extracellular liquid contains all sorts of chemicals, many with electric charges, that influence cellular function. The cell membrane keeps chemicals that are inside and outside the cell separate. At rest (when the cell is not being influenced by inputs), the chemical composition of the inside of the cell is more negatively charged than the fluid outside, due to the kinds of ions that are present on the other side of the cell membrane. Neuroscientists have measured the difference in electrical charge between the inside and outside of a nerve cell. In general, the inside of a neuron that is not being stimulated is about 60 millivolts (60 one-thousandths of a volt) more negative than the outside. In other words, the resting potential of the cell is about -60 mV.

For our purposes, the actual voltage is not that important. All we need to keep in mind is the fact that the membrane potential is fairly negative at rest. When a neuron is stimulated by excitatory inputs from other neurons, however, the membrane potential becomes more positive. The reason for this is related to the way that glutamate works as a neurotransmitter.

Glutamate receptor molecules span the cell membrane, with part facing inside the cell and part facing outside. When glutamate (released from a presynaptic terminal) binds to the outside part of a postsynaptic receptor, a passage opens up through the receptor, allowing positively charged ions in the extracellular fluid to move inside the cell, which changes the chemical balance between outside and inside. If enough glutamate receptors are occupied on the postsynaptic cell at about the same time, and the voltage inside becomes sufficiently positive, then an action potential occurs.

In contrast, when GABA receptors are occupied, the inside of the cell becomes more negative (due to the influx of negative ions, especially chloride, through a passage in the GABA receptor). This makes it harder for glutamate released from other terminals to change the concentration of the positive ions in the postsynaptic cell sufficiently to trigger an action potential. Whether an action potential occurs, then, depends on the relation between glutamate (excitation) and GABA (inhibition). And since any one cell receives many excitatory and inhibitory inputs from many other cells, the likelihood of firing at any one moment depends on the net balance between excitation and inhibition across all of the inputs at that particular time.

Glutamate receptors tend to be located out on the dendrites, especially in the spines, whereas GABA receptors tend to be found on the cell body, or on the part of dendrites close to the cell body. In order for glutamate-mediated excitation to reach the cell body to help trigger an action potential, it has to get past the GABA guard. Excitation coming down a dendrite and headed for the cell body can be extinguished by GABA.

Without GABA inhibition, neurons would send out action potentials continuously under the influence of glutamate, and would eventually literally fire themselves to death. This effect has been demonstrated in experiments where the action of GABA is blocked artificially, or where powerful doses of glutamate-related compounds, too strong to be inhibited by natural levels of GABA, are administered. Overactivity of glutamate, and the resulting injury to neurons, actually plays an important role in stroke and other vascular disorders of the brain, as well as in epilepsy and possibly Alzheimer's disease. Some people have experienced mild versions of glutamate toxicity after eating Chinese food. Monosodium glutamate (MSG), sometimes used as an additive in this cuisine, can increase the amount of glutamate in the body to the point of causing headaches, ringing ears, and other physical symptoms. Regulation of GABA inhibition is one of the ways that psychoactive drugs work. For instance, the anti-anxiety drug Valium works by enhancing GABA's natural ability to regulate glutamate. Excitatory inputs that would normally elicit anxiety by firing action potentials in fear circuits are less able to do so in the presence of Valium and related drugs.

MOD SQUADS

Interactions between glutamate and GABA are key to understanding information processing by the brain, but these substances do not work alone or in isolation. For example, when receptors in the eye detect patterns of light, they send messages through the axons of the optic nerve to the brain. When the electrical signal reaches the axon terminal, glutamate is released. (Whether the postsynaptic cell fires depends not only on the counterbalancing force of GABA inhibition, but also on other chemicals that are present at the time. These are called modulators.

Modulators are neurotransmitters in the sense that they provide a chemical link between the site from which they are released and the location of the

receptors upon which they act. But in contrast to glutamate and GABA, they are less directly involved in the transfer of information from point to point in hierarchical circuits. The way a modulator is distinguished from a transmitter is different for different kinds of modulators, as we'll see soon. And sometimes, the distinction is murky. But one important difference is related to their speed. Glutamate and GABA are fast-acting: they cause an electrical change in the postsynaptic cell within milliseconds of being released from the presynaptic terminal, and their effect is over in a matter of milliseconds. Modulators, on the other hand, have slower and longer-lasting effects.

We'll consider three classes of modulators: peptides, amines, and hormones. Each can have excitatory or inhibitory effects, depending on the specifics of their participation in functional circuits.

Peptides represent a large class of slow-acting modulatory substances found throughout the brain. They are made up of many amino acids, and are larger molecules than simple amino acids like glutamate or GABA. Because peptides are often present in the same axon terminal as glutamate or GABA (but in their own separate storage compartments), they are released with the fast transmitter when an action potential comes down the axon. But peptides bind to distinct postsynaptic receptors and can, as a result, augment or reduce the effect of the fast transmitter with which they are released. However, since peptides are slow to affect the postsynaptic site, and their effects are long-lasting, they tend to have more of an effect on subsequent squirts of fast transmitter. While glutamate and GABA can have slow effects as well as fast ones, depending on the receptors involved, peptides typically only have slow modulatory actions. They can affect dramatically the ability of a cell to be fired by other inputs, but cannot do so with precise timing.

There are many, many peptides that participate in a wide variety of bodily functions. Our interest is in the neuroactive peptides, those that act in the nervous system. The best known of these are the opiates—endorphins and enkephalins. These are triggered by pain and stress and bind to their special receptors, altering pain

sensations and mood. "Jogger's high" is said to be an opiate effect. Morphine generates its effects by binding to these receptors.

The monoamines, another class of modulators, include substances like serotonin, dopamine, epinephrine, and norepinephrine. Unlike most other transmitters and modulators, the cells that produce monoamines are found in only a few areas, mostly in the brain stem, but the axons of these cells extend to widespread areas throughout the brain (figs. 3.8 and 3.12). In this way, a small number of highly localized neurons making monoamines can influence cells in many other locations. Monoamines achieve their effects by facilitating or inhibiting the actions of glutamate or GABA (and the peptides that are released with them). Because the axons are so widely distributed, monoamines have relatively nonspecific effects. They are thus not involved in precise representation of stimuli in specific circuits. Instead, monoamines produce global state changes in many brain areas simultaneously, such as the high degree of arousal occurring throughout the brain when we encounter a sudden danger or the low degree of arousal required when we are going to sleep.

Many drugs used in the treatment of psychiatric disorders work by altering monoamines. Prozac, for example, prevents the removal of serotonin from the synaptic space. Normally, as part of the process by which transmitter action is regulated, neurotransmitters are sucked back into the terminals that release them. By preventing the removal of serotonin, allowing more to stay around longer, Prozac amplifies its effects. One theory holds that there is a deficiency of serotonin in depressed or anxious brains, which Prozac helps correct. The exact means by which the increase in serotonin levels relieves anxiety or depression is not known.

Antidepressant drugs (like monoamine oxidase inhibitors and tricyclic antidepressants) and antipsychotics (like chlorpromazine or phenothiazine) also work by altering monoamine levels. Amphetamines are also targets of recreational drugs: cocaine and amphetamine affect norepinephrine and dopamine levels, while LSD acts on serotonin receptors.

Another monoamine is acetylcholine, which functions as a fast transmitter when it works with one receptor and as a modulator with a different receptor. Disruption of acetylcholine in the neocortex is believed to play a role in Alzheimer's disease, and many drugs that have been tested as treatments for Alzheimer's alter acetylcholine function. Acetylcholine is also a very important transmitter in the body, involved with nerves such as those that control muscle movements and heart rhythm. Nerve gas works by disrupting acetylcholine transmission at muscles, especially muscles required for normal breathing. Many insecticides have similar effects in bugs.

Hormones are the last class of modulators we will consider. Typically, they are released from bodily organs (like the adrenal, pituitary, or sex glands) into the bloodstream where they travel to the brain. There they can, like other modulators, alter the efficacy of glutamate or GABA transmission by binding to specific receptors on cells. For example, cortisol, a steroid hormone released from the adrenal gland during stress, is known to alter information transmission in a variety of circuits involved in memory and emotional processes, in part by altering the ability of GABA to inhibit glutamate. Sex hormones, such as testosterone and estrogen, also can have profound effects on neural transmission and other brain functions. The mood-altering effects of monthly variation in estrogen levels in females are widely discussed, and estrogen replacement therapy during and after menopause is believed to counter some of the effects of aging on brain functions. Because hormones reach the brain through the bloodstream, they can influence many regions simultaneously. However, since only certain areas, and only certain circuits in those areas, possess the relevant receptors, considerable specificity can be achieved by hormonal modulation.

GOLGI AND THE GAP

As important as chemical synaptic transmission is in the brain, another form, called electrical transmission, also occurs. Although the extent to which electrical synapses operate is not known, it is becoming more and more apparent that they are significant forces for us to deal with as we conceive brain function.

In order for two neurons to communicate electrically, their membranes have to fuse in such a way as to allow the direct flow of electricity from one to the other. These points of fusion are called gap junctions. Recent studies have shown that in some brain areas, like the hippocampus a region important for the formation of explicit memories, GABA (inhibitory) cells are linked together or electrically coupled, by gap junctions. In this way, when GABA cells are activated, excitation can spread between them in such a way as to activate many of the interconnected cells at once. The cells then fire together, in synchrony, and thereby can regulate activity of projection cells throughout the region.

The existence of gap junctions gives partial vindication to Golgi's reticular theory of the brain, in the sense that some neurons can communicate directly by way of physical fusion. Much remains to be learned about them, and their contribution to synaptic transmission needs to be better integrated with our knowledge of chemical transmission.

CIRCUITS IN ACTION

The same basic transmitters, modulators, and hormones can be involved in very different functions. Our abilities to see, hear, remember, fear danger, and desire happiness all involve excitatory (glutamate) synaptic transmission regulated by inhibitory (GABA) synapses and modulated by peptides, amines, and hormones. What makes a sound different from a sight, a memory different from a perception, a fear different from a desire is not so much the chemistry involved but instead the specific circuits in which the chemicals act. As a way of illustrating how glutamate, GABA, and modulators work, let us consider their role in the detection of danger by the amygdala.

The amygdala detects danger by virtue of its position in a synaptically connected system. In its simplest form, this system can be described in terms of a three-level excitatory chain of cells that releases glutamate—projection cells in sensory systems activate projection cells in the amygdala, which activate projection cells in motor control areas. This scheme leaves much out, but we'll have the opportunity to embellish it later.

Amygdala cells receive inputs from the sensory world constantly, but they ignore the majority of them. In fact, they tend to be quiescent most of the time. They do get worked up, though, when the right kind of stimulus is present—one that signifies danger or some other biologically significant event. This has been shown to be true in studies of both lower animals and humans. So what keeps a projection cell in the amygdala from firing in response to meaningless stimuli? The answer, as you've probably guessed, is GABA.

As we've seen, the resting membrane potential of cells in many brain areas is about -60mV . In the amygdala, however, some cells can be as negative as 80mV , due to sustained or tonic inhibition by GABA. With GABA receptors on amygdala projection cells occupied and passing chloride, the inside of the cells becomes more negative, which means it takes extra excitation to turn the amygdala on. As a result, not any old stimulus will do the trick. The stimulus has to have special qualities that allow it to overcome the tonic inhibition produced by GABA.

Stimuli that are inherently dangerous (the sight or smell of a predator) or unpleasant (intense stimuli, like loud noises or stimuli that cause pain) are able to overcome the tonic inhibition, as are stimuli that have emotional resonance acquired through past learning. Thus, an otherwise meaningless sound of modest intensity that previously occurred in association with pain has the same effect as a natural (innate) form of danger. Both innate (hard-wired) and learned danger signals cause amygdala cells to fire rapidly for a sustained period, and are thus able to overcome the GABA guard.

Even after fear-arousing stimuli get past tonic inhibition and cause amygdala cells to fire, however, they are still subject to GABA control. The inputs to the amygdala activate GABA cells as well as projection neurons. As a result, as the inputs become more active, the elicited inhibition in the amygdala builds up, which in turn begins to shut down the activity of amygdala cells.

If the ability of GABA to keep meaningless stimuli from turning on the amygdala is compromised for some reason (either because the projection cells come to fire more easily or because the GABA cells fire less easily), stimuli that are not dangerous come to be responded to as though they were. This may occur in certain fear and anxiety

disorders. By the same logic, things that make projection cells fire less readily or that make GABA cells fire more readily should reduce fear and anxiety. Indeed, one of the most popular medications ever invented for the treatment of anxiety is Valium, which works by facilitating GABA transmission. Although drugs taken orally reach many sites in the brain, it is likely that at least some of their effects on fear and anxiety are achieved by enhancing inhibition in the amygdala, and thereby making it harder for external or internal stimuli to elicit fear responses by activating amygdala circuits.

The amygdala also receives modulatory inputs of various types. For example, serotonin fibers terminate there, and when the amount of serotonin rises in the amygdala the activity of excitatory projection cells is inhibited. The inhibition in this case is not due to the fact that serotonin directly affects projection cells, but rather that serotonin excites GABA cells, and thus increases the degree to which they inhibit projection neurons.

Drugs like Prozac work by increasing the amount of serotonin available at synapses. By enhancing serotonin transmission at GABA synapses in the amygdala, and thereby reducing the activity of projection neurons, Prozac may, like Valium, help control anxiety by reducing the ability of inputs to the amygdala to activate fear circuits.

The amygdala is also the target of many hormones. One of these is cortisol, which is released from the adrenal cortex during fear-arousing and other-wise stressful events. The facilitation of GABA inhibition of amygdala projection cells by serotonin is modulated by cortisol. Serotonin's ability to facilitate inhibition by exciting GABA cells thus depends on the binding of cortisol to receptors located on amygdala neurons. Cortisol is elevated in a variety of psychiatric disorders, and cortisol increases the intensity of fear reactions. Drugs like Prozac may reduce exaggerated fear and anxiety in psychiatric disorders by enhancing the ability of serotonin to facilitate GABA inhibition in the presence of elevated cortisol.

The fear system thus nicely illustrates the basic elements of neural transmission in the brain and its regulation by modulatory chemicals. We will build upon these points at various times in later chapters.

ARE SYNAPSES ENOUGH?

My emphasis on the importance of synapses in brain function is not intended to minimize the role of other factors. For example, the rate at which a cell fires spontaneously is a function of certain electrical and chemical characteristics of the cell. These are called intrinsic properties to distinguish them from extrinsic influences from other cells mediated by synaptic transmission and modulation. A cell's intrinsic properties, which may have a strong genetic component, will greatly influence everything that cell does, including its participation in synaptic transmission. But because psychological and behavioral functions are mediated by aggregates of cells joined by synapses and working together rather than by individual neurons in isolation, the contribution of the intrinsic properties of a cell to mental life or behavior occurs only by way of the role of that cell in circuits. While synapses themselves don't account for everything the brain does, they do participate crucially in every act or thought that we have, and in every emotion we express and experience. Synapses are ultimately the key to the brain's many functions, and thus to the self.

CHAPTER FOUR

BUILDING THE BRAIN

EACH CHILD IS AN ADVENTURE INTO A BETTER LIFE—AN OPPORTUNITY TO CHANGE THE OLD PATTERN AND MAKE IT NEW.

—Hubert H. Humphrey

The brain's billions of neurons are intricately connected in ways that make possible the mundane (such as the regulation of breathing) and the marvelous (the belief in an idea). But how do cells in the developing embryo become neurons, and how do they end up in just the right places? How do the axons of all these cells find their way to their target areas? And once having reached them, how do the terminals figure out exactly which neurons to make synapses with? Because the various steps take time, and because different circuits go through these steps on different schedules, our behavioral and mental repertoire unfolds gradually, and unevenly, during childhood. Somehow, though, it all comes together, and a person, a self with all its aspects, emerges.

Brain development is the major battlefield of the nature-nurture conflict. In its simplest form, the debate is about whether mental and behavioral characteristics are determined more by genes or by environment. To the extent that mental and behavioral characteristics are functions of the brain, and synaptically connected circuits underlie brain functions, the nature-nurture debate essentially reduces to questions about how circuits are built during development.

No one today seriously proposes that the brain is a blank slate at birth, waiting to be written on by experience, or, alternatively, that it is a genetically predetermined, unchangeable repository of tendencies to act, think, and feel a particular way. Instead, it is widely believed that brain circuits come about through a combination of genetic and nongenetic influences, and the debate now hinges less on the dichotomy than on the manner in which nature and nurture contribute to brain construction.

The dichotomy, in fact, begins to dissolve when it is realized that, regarding questions of mind and behavior, nature and nurture are really two ways of doing the same thing—wiring up synapses—and both are needed to get the job done. We commonly think of experiences as leaving their mark on the brain through the record of memory, and, as we'll see in later chapters, memory is a product of synapses. It is less common, but no less appropriate, to think of genes as also influencing us in the form of memory. In this case, though, the synaptic memory comes about as a result of ancestral rather than personal history. The shaping of synaptic connections in early life by genes and experience is the topic of this chapter.

FIRST STEPS

The earliest events in embryonic brain development are largely controlled by genes, their products and the local chemical environment in which they exist. The job of genes is to make proteins, which regulate many aspects of brain development. Some proteins are enzymes that trigger chemical reactions, others induce additional genes to

make additional proteins, some form barriers that guide and restrict the many cell movements that take place, and still others provide adhesive surfaces on cells to which other cells cling while making their way to their final destinations. When we speak of genetic influences on brain development, we are essentially describing the effects of proteins and their chemical spin-offs.

The young embryo, lacking sensory systems, is largely isolated from direct perceptual contact with the external environment. But even in the earliest stages of development, genes do not operate completely independently of the outside world. The chemical environment of the embryo is, by necessity, in direct contact with the body chemistry of the mother. The embryo cannot make on its own the amino acids that are used to assemble the proteins that are required for brain and body development. They have to be obtained from the mother, who gets them from the food she eats. The mother's diet can also be the source of less desirable substances—toxins and chemical additives in foods, for example—as can the air she breathes, drugs she takes (prescription

and recreational), and cigarettes she may smoke. The mother's level of stress will affect her hormonal state, which can influence the embryo, as can antibodies she makes to fight off infections. Although the major features of the brain are dictated by a genetic plan (which guarantees that all human brains look and work much the same), this plan requires certain conditions in the internal chemical environment in which neurons will grow. If this gene-environment interaction is perturbed, so, too, will be the normal development of the brain. Nature and nurture interact from the start.

Brain development begins in the ectoderm, which, together with the mesoderm and endoderm, make up the three major parts of the embryo. These give rise to the various regions and organs of the body. A thickening of a part of the ectoderm forms the neural plate, which folds to form the neural tube, out of which the brain begins to be constructed.

In the neural tube, nonneural precursor cells divide and give rise to neurons. In humans, the vast majority of neurons are made in the months just prior to birth. At the peak production point, about 250,000 neurons are being generated per minute. This process is controlled by hormones that diffuse up into the neural tube from underlying tissues and turn on genes that make proteins, which, in turn, regulate neuron production by precursor cells. (The hormones themselves are gene products made in nonneural parts of the embryo.) If production of neurons is interfered with during this early period (say, by consumption of alcohol or drugs), birth defects, such as spina bifida or anencephaly, can result. However, not all birth defects are due to drug or alcohol consumption, and drug or alcohol consumption does not always lead to birth defects.

According to conventional wisdom, no new brain cells are formed in adult mammals—the brain cells you acquire in early life are pretty much the ones you die with, minus the ones you lose for various reasons over the course of years. However, recent studies by Fernando Nottebohm and Bruce McEwen of Rockefeller University, Elizabeth Gould of Princeton, and others have challenged the conventional view, demonstrating that neurons continue to

generated in some areas of the adult brain.) At this point, it is not known how widespread the birth of new neurons is, and it may be fairly rare. However, this finding offers at least the glimmer of a hope that brain scientists might someday be able to develop therapies fostering the growth of new neurons for largely untreatable neurological disorders, such as spinal cord injuries, stroke, and Alzheimer's disease. It's important to make a distinction here between neurogenesis—the birth of new neurons—and synaptogenesis—the creation of new synapses between existing neurons. Synaptogenesis is a common phenomenon, one that probably occurs up until the moment of death. As we will see in this chapter, synapses change dramatically during early life, and in the next two chapters, we'll see how synapses are changed every time our brain records an experience.

Soon after their birth, neurons begin to segregate, with cells destined to be part of the hindbrain, midbrain, and forebrain taking up different territories in the growing neural tube. Segregation is under the direct control of a special set of genes called homeotic genes, which make proteins that control the placement of cells, providing boundaries that guide and restrict cell movement, and adhesive surfaces that allow cells to cluster together. Some scientists

believe that autism might be due to a mutation of homeotic genes that leads to faulty brain construction and connection) The function of homeotic genes was discovered in studies of fruit flies, and is being elaborated on in studies of flies and worms, especially the worm *C. elegans*. It may be difficult to believe that we have much in common with flies and worms, but the fact is that these genes, as well as a number of others, have been preserved through many levels of evolutionary history.

Segregated cells eventually differentiate, which means that they take on different shapes and sizes, and end up with different transmitters and modulators. This is how, for example, projection and interneurons come to differ. Differentiation is under genetic control but is not strictly a genetically determined process. Studies have shown that if cells from one brain region are transplanted to a different region in another animal prior to differentiation, they take on the characteristics expected of cells in the region of the host rather than that of the donor. Chemical factors in the local environment appear to determine the ultimate type of cell that will be expressed. However, the switch can only occur in very young cells, indicating that once a cell's type is determined, its fate is sealed. This suggests that cell type is not rigidly dictated by genes and is strongly influenced by the environment, albeit in this case the local chemical environment that surrounds the cell rather than the environment that is external to the organism as a whole. The local cues involved, though, are proteins that have themselves been genetically coded, so the nongenetic contributions to cell differentiation are not too far removed from genetic ones.

As development proceeds, the neural tube expands and bends, eventually assuming a shape that begins to resemble the brain. While segregation initially places cells destined to be in the same brain regions together in the neural tube, as the tube grows, the cells have to migrate out from their segregated

resting place to reach their final destinations in the growing brain. Migration, for example, allows the many regions that make up the adult neocortex to receive from the neural tube the correct numbers and types of cells that will be needed for those areas to perform their functions.

Although the exact manner in which migrating cells find their destinations is not completely understood, work on the development of the cortex by Pasko Rakic of Yale suggests that the process involves the building of scaffolds or chemical trails that migrating cells follow. These trails are formed by glial cells (nonneural cells found in the brain that contribute in important ways to development and other aspects of brain function). The glia are themselves guided by local chemical cues, made by genes and their by-products, that serve as molecular signposts, creating barriers that restrict movement and providing adhesive surfaces. As the emerging brain expands in size, the traveling glia grow fibers that extend out from the proliferative zone (the place where the precursor cells are making neurons and glia) toward the brain surface (where the neurons need to go to form the cortex). By crawling along the glial trail, the young neurons find their way to their target.

The glial trail makes it possible for some of the neurons living next to each other in the segregated neural tube to end up next to each other in the cortex. This accounts in part for how cortical areas come to have the cells they do—those adjacent to each other in the tube travel together along the same glial fibers and tend to end up in the same cortical location. However, the manner in which cortical or other brain areas are assembled is not fully understood. One possibility is that the function of a cell (whether it will be involved in sight vs. touch sensations, for example) is specified genetically, so that it is destined to end up in specific areas devoted to a specific function before it leaves the neural tube. Another possibility is that cells learn what to do from the context in which they end up. In support of the second possibility is the fact that if a piece of visual cortex is removed and transplanted into the location of the somatosensory (touch) cortex, it will accept somatosensory axons from a lower station in the brain (the somatosensory thalamus) and then function as somatosensory cortex.⁷ This suggests that cells adapt to the local situation (the chemical environment and the axons that are coming into it) in the process of taking on their function. Support for the first possibility (that the functional fate of a cell is genetically determined) comes from a recent study showing that under certain conditions where axons are prevented from growing into the cortex, the cells in the cortex still come to possess the chemical molecules that characterize that region. This suggests that the chemical signature of the cell is determined by the cell itself, likely by its genes. However, from the point of view of brain function, the job

performed by a piece of cortex is often more dependent on its synaptic connections than on the molecules present. As a result, even if the chemical signature of a region is fixed by genes, the function of the region will in all likelihood be more determined by its synaptic connections. Cells traveling along the glial trail thus seem to be ambivalent functionally, at least to some extent.

Once neurons reach their destinations, they begin to sprout axons. These fledgling fibers then have to find their way to their targets, the neurons with which they will form synapses. The targets can be located nearby or in faraway areas, or both. Their so-called pathfinding depends on

growth cones, structures shaped like Japanese fans that extend from the ends of growing axons. The tips of the growth cones have chemical affinities with specific proteins in the surrounding space—they only attach to certain proteins, and are repelled by others. When one of the tips of a growth cone finds the right substance, it attaches. If the other tips are not attached, the growth cone is twisted in the direction of the attachment, and the axon follows. Recent studies have shown that the same protein can be both an attractant and a repellent for an axon, depending on the chemical state of the axon.⁰ Slowly, step by step, the growth cone crawls along the chemical trail, pulling the axon forward, until it reaches its destination. Some signal, the nature of which is still not fully understood, then causes the growth to stop. The first axon along the trail, the pioneer, leads the way, making a path that other nerve fibers from the same area can follow. Pathfinding by fibers is not unlike migration along the glial trail by cells.

When axons reach their target, and receive the signal to stop growing, growth cones give way to axonal terminals, which begin to form synaptic connections with postsynaptic cells. The initial connections, like the other early events in brain development already discussed, are largely established by genetic and other intrinsic factors. But the brain is not completely built from the inside out. The final steps in its construction, especially the steps through which connections are fine-tuned in a way that enables them to do their jobs, require more.

PRO-CHOICE

The human genome contains instructions for body development, including, as we've seen, the development of the overall structure of the brain. Somewhere between 50 and 70 percent of all genes in the human body are believed to be in the brain. Nevertheless, we probably don't have enough genes to account for the wiring of each of the brain's trillions of synaptic connections."

Pretty much everyone agrees that the transition from the uncommitted, immature initial connections of the young brain to the mature and very specific connectivity that characterizes that of the adult requires neural activity, that is, transmission across synapses. The exact role of neural activity, though, is heatedly debated, with the main issue of contention being whether activity, especially activity initiated by environmental stimulation, helps create the mature connections or just selects from the initial set of intrinsically established connections those that will be retained. This instruction vs. selection debate cuts deep into the heart of human nature: Is the self sculpted from a preexisting set of synaptic choices, or does experience instruct and add to the synaptic basis of the self as we go through early life? As we will see, since environmentally triggered neural activity is involved in both instructing and selecting connectivity, this is not so much a debate about genes vs. environmental experience as one about the precise contribution of experience.

Selectionist ideas originated in evolutionary (Darwinian) biology, were adopted and adapted by the field of immunology, and were then applied to brain function. Niels Jerne, a Nobel Prize recipient in immunology; introduced selectionist thinking to the study of the brain in the late 1960s. Jerne pointed out that the history of biology is filled with instances of instructional ideas giving way to selectionist ones. In his own field of immunology, for example, it was once thought that foreign antigens enter cells and instruct them to make antibody molecules. That model is now known to be incorrect. To do this, a cell would have to be able to identify foreign antigens. Since the recognition of antigens is the business of antibodies, the only way to make an an-

tibody to a novel antigen would be to already have the antibody. Jerne saw the fallacy in this logic, and subsequent research showed, as he predicted, that foreign antigens select precursor molecules from a preexisting

pool that can be assembled into a large variety of antibodies.

Jerne later applied his antibody logic to the topic of learning. He objected to the instructional implications of John Locke's seventeenth-century empiricist hypothesis that the mind is a blank slate filled in by experience, and instead sided with the Greek Sophists, who believed that learning, as such, is impossible. He suggested that the idea of learning from experience (instruction) be replaced with the concept that experience just selects from preexisting latent knowledge. Paraphrasing Socrates, Jerne noted that "learning consists of being reminded of what is already in the brain."

A few years after publication of Jerne's hypothesis, Jean-Pierre Changeux, a prominent French neuroscientist, used selectionist principles to explain certain findings he had obtained regarding the role of neural activity in the development of synaptic connections between nerves and muscles. He concluded that neural activity "does not create novel connections but, rather, contributes to the elimination of pre-existing ones.

Perhaps the most vocal contemporary practitioner of neural selectionism is Gerald Edelman, who like Jerne received a Nobel Prize for his work on the immune system. In *Neural Darwinism*, Edelman argued that synapses in the brain, like animals in their environments, compete to stay alive. Synapses that are used compete successfully and survive, while those that are not used perish. According to Edelman, "The pattern of neural circuitry... is neither established nor rearranged instructively in response to external influences. External influences, instead, select synapses by initiating and reinforcing certain patterns of neural activity that involve them.

Selectionists assume that genetic and nongenetic factors interact at each step of brain development. Selection operates on preexisting connections set up by genes (which make proteins that help guide axons to the right areas working in concert with nongenetic factors (chemicals from the mother, for example). But genes and the chemical environment are not wholly responsible for establishing the initial connections. Selectionists also assume that there's a good deal of randomness involved—terminals and dendrites that happen to be in the same vicinity take the opportunity to form synaptic connections, independent of the overall guidance plan specified by genes. As a

sult, in spite of a general genetically programmed plan, the preexisting connections upon which selection ultimately operates also have a unique, individualistic nature, from which experience then does the selecting. Because each person's experiences are different, different patterns of connectivity are selected. Genes thus dictate that we will all have a human kind of brain with roughly the same kinds of circuits, but random individual differences will exist, and the connectivity of the circuits, selected by synaptic activity, will shape the individual brain.

SYNAPTIC MATH

Three of the main tenets of neural selectionism are described by the terms *exuberance* (more synapses are made than are kept), *use* (the synapses that are kept are the ones that are active), and *subtraction* (connections not used are eliminated). Building the nervous system in this fashion is thought to provide a means of coping with the paucity of information available to the brain from other sources, such as the outside world, during early development.⁷ If the selectionists are right, the connections we end up with as adults are those that were not subtracted—the self, they would say, is not constructed, it is selected from preexisting possibilities.

Pioneering studies by Viktor Hamburger and Rita Levi-Montalcini in the 1930s showed that subtractive or regressive events do occur in the nervous system, a view that is now widely accepted. During development, cells are made and then are killed off; chemicals come and go; functions change. Our main concern here is with synaptic regression, the pruning back of exuberant, unused projections during early development.

Evidence for synaptic regression was provided in a widely cited study by Pasko Rakic and his colleagues who found that the number of synapses in several areas of the primate cortex rises and then falls during the first year of life. However, a follow-up study involving a more detailed analysis of one cortical area indicated that the decrease

didn't occur until puberty. Given that cognitive development is close to completion by puberty, a decrease in synapses that doesn't occur until puberty can't account for the maturation of the mind. Although few such investigations have been performed in humans, one widely cited study suggests that across the cortex the greatest number of synapses are present at around twenty-four months of age (obviously long before puberty), but with different areas reaching their peak at different times.

Certainly, much more work is needed to determine the extent to which synapse elimination occurs, when it occurs in specific neural circuits, whether there are significant species differences, and how such data should be interpreted. The key issue is not whether synapses are eliminated during development (they clearly are) but instead whether the decreases that do occur provide conclusive support for the strong version of selectionism—the version that claims that activity *only* prevents the elimination of synapses. This takes us to the second part—the “use it or lose it” part—of the selectionist argument.

USE IT OR LOSE IT

That use prevents synaptic demise was suggested by a set of classic studies started in the 1960s by David Hubel and Torsten Wiesel on visual system development.⁷ Like most other sensory systems, the visual system uses receptors near the body surface (in this case, cells in the part of the eye called the retina) to take in information from the environment. Retinal cells give rise to axons that terminate in several postsynaptic targets, the major one of interest here being the visual thalamus. Thalamic cells, in turn, send their axons to areas of the visual cortex, where, to state it most simply, our visual perceptions emerge.

Neocortical areas, by definition, have six layers. In the visual cortex, the axons from the thalamus end mostly in the middle layers and are distributed to the rest of the visual cortex from there. (Axons from regions other than the thalamus end in other layers.) In studies of mature cats, Hubel and Wiesel found that cells in these middle layers tended mainly to fire action potentials when stimulated visually through one eye or the other, with fewer cells responding to stimulation of both eyes. The eye-specific cells tended to cluster together in separate patches, leading to areas with cells responsive to only one eye. But when they performed the same study on very young cats, most cells responded to stimulation of both eyes. This led them to surmise that initially, cortical cells must receive synaptic inputs from each eye, and then, during later development, inputs from one eye were eliminated, leaving only the inputs from the other.

To test whether their hunch was correct, Hubel and Wiesel closed one eye of their experimental subjects early in life before the eye segregation took place in the visual cortex. The eye was opened when the animals became adults. By then, relatively few if any cells responded to visual stimulation of the previously closed eye. Nevertheless, the cells that would have normally responded to the closed eye were still functional—but now they only responded to the eye that had been left open. The open eye, in other words, retained access not only to its own cells but also to the cells normally devoted to the other eye. Closure of an eye in an adult cat had no such effects, suggesting that the cortex only has this capacity to reorganize during early development. Hubel and Wiesel concluded that the development of synaptic connections is a competitive process, one in which the connections that are used are kept and those that go unused are eliminated.

These findings, which have been confirmed and elaborated on by many investigators over the years,²⁹ are certainly consistent with the selectionist principle that adult synapses are picked from a preexisting set by activity. However, since this early study was conducted, the tools available for exploring how the brain works have advanced considerably, allowing neuroscientists to now ask more sophisticated questions. Michael Stryker and his colleagues, for example, studied the issue using a more exacting measure of the effects of

synaptic activity.” Rather than measuring the overall clustering of eye-specific cells in a region, they asked what happens to individual axon terminals and the synapses they make during development. Individual cells in the thalamus were injected with a chemical tracer at different times.” The tracer then spread throughout the neuron, all the way to the axon terminal and its presynaptic endings in the cortex. They found that as development proceeded,

many axons did indeed retract (consistent with a selectionist view), but that the remaining axons also increased in complexity (consistent with an instructional view). Activity thus leads to an increase in synaptic complexity rather than merely a stabilization of the preexisting pattern—activity therefore is capable of instructing the formation of new synaptic connections.

Activity also appears to help define the demarcation between areas of the cortex. For example, the axons from the visual thalamus normally spread into areas of the auditory cortex, and vice versa, in early life. As development proceeds, the stray connections are pruned back, since the cells in a given area are more strongly connected with one system or the other, and the stronger inputs compete more successfully for the attention of the postsynaptic cells. This can be seen with congenitally deaf adults, in whom the region of the cortex responsive to visual stimulation includes the area that would typically be occupied by the auditory cortex. It is the absence of activity in the auditory area that has allowed the visual fibers to compete for those synapses.

While we would tend to assume that synaptic activity in a sensory system, like the visual system, is due exclusively to environmental stimulation, studies by Carla Shatz and others suggest that a broader view of such activity is needed. In contrast to cats and some other animals, primates have eye-specific patches present in the cortex at birth. Several weeks before birth, cells in the retina are spontaneously firing action potentials in these animals. Because many cells in a given eye fire simultaneously, a wave of synchronized activity is transmitted to the brain from each eye. However, the waves coming out of the two eyes are not in sync and, in fact, typically occur at different times. This suggests that endogenous activity generated in the two eyes independent of external visual stimulation might be sufficient to provide the competitive activity in the cortex to create eye-specific clusters of cells. Indeed, if this endogenous activity is blocked chemically, the clusters fail to develop.

It's important to make a clarification at this point. New connections formed by activity are not created as entirely new entities but rather are added to intrinsically determined preexisting connections. Added connections are therefore more like new buds on a branch rather than new branches. Activity thus does not produce wholesale rewiring of the brain. After all, most of the connections in my brain and yours are the same. Activity adds those little adjustments that make you and me different.

In reviewing the evidence on the role of activity in the establishment of the adult pattern of synapses, two leaders in the field, Larry Katz and Carla Shatz, recently noted that “neural activity is likely providing cues that drive the formation of new synapses and axon branches, as well as cues that act to select and stabilize existing ones.” While it may seem that such an “everybody wins” summary statement implies that nothing has been accomplished in the field, this would be the wrong conclusion. It has taken a tremendous amount of work to demonstrate that activity both instructs the formation of new connections and selects used ones from a preexisting pool. While activity may not be necessary for the establishment of all circuits, the fact that it plays a key role in constructing at least some of them suggests that a purely selectionist view of brain development is not viable. At the same time, the fact that activity instructs the growth of connections does not rule out its role in selecting and eliminating synapses. Activity is involved in both processes.

It is probably best to think of instruction and selection as two complementary means by which circuits can be constructed rather than as mutually exclusive theories of brain development. Unfortunately, the only way to determine the precise contribution of each to the construction of a particular circuit is to do the experiments. Since only a few circuits have been studied in this degree of detail, developmental brain researchers are not likely to be bored any time soon.

WIRING BY FIRING

New connections made by neural activity are just as subject to the “use it or lose it” rule as old ones that have been selected from the initial pool of synapses. When new connections are made, only some survive—those that go on to be used. But how exactly does “use” prevent loss? What occurs between a postsynaptic cell and its presynaptic terminal to ensure a connection's survival? Or, stated more positively, how does activity strengthen

connections?

In 1949, Donald Hebb, a Canadian psychologist, proposed that if two neurons are active at the same time, and one is presynaptic to the other, then the connection between them will be strengthened. In Hebb's words, "When an axon of cell A is near enough to excite cell B or repeatedly and consistently takes part in firing it, some growth process or metabolic changes take place in one or both cells such that A's efficiency, as one of the cells firing B, is increased." The essence of Hebb's idea is captured by the slogan "Cells that fire together wire together."

Although Hebb originally proposed his fire-and-wire theory to account for the nature of learning and memory, it has been used to explain other aspects of synaptic function, especially the construction of synapses during development. Consider again the establishment of visual cortex connectivity. As we've seen, in primates, the weeks before birth are an important period for visual system development, as waves of spontaneous activity from the two eyes set up patterns of activity that result in the preferential activation of certain cortical cells by one eye or the other. Because cells in the retina of one eye are more likely to fire spontaneously at the same time, and are much less likely to fire at the same time as cells in the other eye, chances are that when a postsynaptic cortical cell is activated by presynaptic inputs from one eye, presynaptic inputs from other cells in the same eye will be arriving more or less simultaneously. According to Hebb's rule, this concurrent activity in presynaptic and postsynaptic cells then leads to a strengthening of the connections from that eye to the postsynaptic cell.

Quite a lot has been discovered about the neurotransmitter mechanisms that underlie Hebbian plasticity; especially in the context of learning and memory. We will examine this work in detail in chapter 6. For now, we can say that the best-understood form of Hebbian plasticity involves the release of glutamate from presynaptic terminals, and the binding of the released transmitter to two kinds of postsynaptic receptors. One of these receptors records that the postsynaptic cell was active, and the other records which presynaptic terminals contacting the postsynaptic cell were active at the same time. In this way, the postsynaptic cell is able to detect the co-occurrence of presynaptic and postsynaptic activity, and thereby help link the postsynaptic cell with all of the presynaptic terminals that were active when the postsynaptic cell fired an action potential.

Hebb suggested that in order for the increased efficiency of the connection between the presynaptic and postsynaptic cell to be maintained, something like new growth needs to occur. Although there was little evidence of it at the time, activity-induced growth has now been observed in many different circumstances. For example, Stryker's results described earlier showed that active axons branch and sprout new connections. Also, studies of Eric Kandel and colleagues at Columbia have demonstrated that physiological plasticity is accompanied by axon branching and new synapse formation both during development and following learning. Once this occurs, a given action potential coming down the axon will be more effective in firing the postsynaptic cell because it activates more synapses on that cell.

The manner in which activity promotes growth during early development is being actively researched. One of the key players is a class of postsynaptic receptors for the excitatory transmitter glutamate (see chapter 3). These special receptors, known as NMDA receptors, appear to be particularly important in that they are able to detect the match between activity in the presynaptic and postsynaptic neuron. When these receptors are blocked or otherwise interfered with, normal development is disrupted. We will have much to say about the role of NMDA receptors in plasticity in chapter 6. For now, it is mainly important to know that mechanisms exist that can indeed allow Hebb's magic to be worked.

Another important set of molecules are neurotrophins, special tonics that promote the survival and growth of neurons. When an action potential occurs in a postsynaptic cell, neurotrophins are released from the cell and diffuse backward across the synapse, where they are taken up by presynaptic terminals. Under the influence of neurotrophins, the terminals begin to branch and sprout new synaptic connections. Since only those presynaptic cells that were just active (that just released transmitter) take up the molecules, only they sprout new connections. Activity thus induces growth, and the growth that occurs is restricted to the active terminals.

In addition to this role in the active construction of circuits, neurotrophins are also involved in synapse selection. The natural fate of many cells during development is an early exit. So-called programmed cell death is one of the regressive events that help shape the final pattern of connectivity. Cell death is prevented if a presynaptic

terminal receives a life-sustaining shot of neurotrophins from its postsynaptic partner. The survival rate of neurons is in this way regulated by the limited availability of neurotrophins. Only those cells that compete successfully for neurotrophins (those that are active) survive. In the presence of neurotrophins, the surviving terminals (those that were active) also begin to sprout new connections. Selection can be a step along the path toward activity-instructed growth—in other words, selection and instruction are partners in synaptic development.

WHAT ABOUT INNATENESS?

The two main ways circuits are known to get built—selection and instruction—both assume that the synaptic connections that ultimately determine what a circuit will do and how it will do it are wired up epigenetically, that is, by interactions between genes and environment (including both the internal and external environment). Where does this leave innateness, the idea that certain functions mediated by the brain might develop mainly under genetic control?

The topic of innateness often comes up in the context of psychological (mental and behavioral) development. To the extent that psychological characteristics are mediated by synapses in the brain, the psychological and synaptic levels are intimately related. The question raised in the preceding paragraph can therefore be rephrased: Where does the epigenetic nature of synaptic development leave the synaptic basis of innateness? In order to consider this question meaningfully, we need to make a distinction between the two different meanings of innateness. One is the idea that genes make us all the same, giving us traits that we share with other members of our species. The other is that genes make us unique.

INNATENESS I: SPECIES SPECIFICITY

For a long time, innateness was the business of ethologists, biologists like Konrad Lorenz and Niko Tinbergen who studied behavior.” A wide range of behaviors in a variety of different animals—from nest-building, maternal care, and foraging to predation and defense—were said to be innate and thus were believed to emerge during early life independent of learning and other environmental influences. Learning was deemed to be important in shaping the individual, but many of the characteristics that make members of a species act alike were thought to be free from environmental influence, or at least more dependent on nature than nurture.

The classic method used by ethologists to test innateness was to show that a behavior was present at birth or shortly thereafter, or, more generally, that it emerged in the absence of an opportunity for the organism to have learned from environmental stimulation. But in the 1950s, Daniel Lehrman, an American animal psychologist, convincingly argued that even under the most stringent isolation conditions, the organism is never completely isolated from nongenetic influences, such as things that happened prenatally or shortly af

ter birth. As a result, the epigenetic nature of behavior development came to be accepted by ethologists, who shifted away from terms like *innate* and *instinctual* and toward *species-typical* in their efforts to understand the development of the unique behavioral repertoires of different species.”

Ironically, just as ethology was coming to accept epigenetic contributions to species-typical behavior, psychology (especially American psychology) turned toward an emphasis on innateness. Today, the innateness of certain mental capacities in humans and other species is emphasized by many psychologists. Below, we'll examine where this trend originated and where it is going.

THE DEMISE OF UNIVERSAL LEARNING.

Through the mid-twentieth century, American psychology paid little attention to special capacities of individual species. Under the influence of behaviorists like John Watson, B. F. Skinner, and Clark Hull, behavior came to be understood as dependent on what organisms learned from their experiences. For behaviorists, learning was a

universal capacity that worked more or less the same regardless of which animal was doing the learning and what was being learned.” If one wanted to understand how humans learn language, math, telephone numbers, tennis, or what someone looks like, studies of the way that key-pecking in a pigeon or bar-pressing in a rat was reinforced by food were just as valid as anything else.

Things began to change in the 1950s, when the linguist Noam Chomsky challenged the notion of a universal learning capacity. He proposed that natural language is unique to humans, and that language acquisition is different from other learning capacities that humans have. His ideas about rules and representations in our use of language helped spark the cognitive revolution in psychology, and his conception of a universal grammar encoded in the human genome has ever since been a cornerstone of cognitive nativism, which assumes that certain psychological capacities are innate.

The plain vanilla view of learning was also challenged when animal psychologists in the 1960s began to realize that the behaviorists’ “law of learning” didn’t apply to all forms of learning. For example, John Garcia’s discovery that some stimuli (those that make you sick to your stomach) can be associated with tastes, but not with sounds or sights, was inconsistent with the behaviorist notion that any given stimulus could be associated with any other one, though it makes perfect sense from a biological point of view. This and other findings suggested that there might be biological constraints on learning, factors that tailor the workings of the learning process to the specific learning task and to the particular learning requirements of the species. A universal learning mechanism was clearly not sufficient, and the study of animal behavior by American psychologists became more ethological.

THE RISE OF EVOLUTIONARY NATIVISM.

Eventually, biological constraints came to play an important role in ideas not just about animal behavior but also about human thought. For example, the psychologist Steven Pinker has recently blended the Chomsky nativist tradition, which views language as a special capacity of humans, with ideas about evolution, asserting that humans have a language instinct, as well as a variety of other innate mental organs (capacities). These are believed to have resulted from Darwinian natural selection and “owe their basic design to our genetic program. That language may be innate is supported by studies of pidgin language, a simple, crude form of verbal communication among people who speak different languages but live together and have to communicate with one another. Pidgin lacks grammar when it is initially pieced together, but by the second generation, it begins to take on a grammatical structure, suggesting that the underlying innateness of grammar ultimately emerges. The innateness of language, though not universally agreed upon, is supported additionally by certain genetic conditions that have been linked to specific chromosomes, including the spared language abilities in spite of other mental difficulties in children with Williams syndrome, and the language impairment in spite of otherwise intact mental abilities in specific language impairment (SLI). Although the evidence for innateness is strongest for language and emotional expressions, innateness has been postulated to account for many mental capacities. For example, it has been proposed that humans have innate knowledge about numbers, physics, and even about how other human minds work.⁷⁰

The recent groundswell of interest in the natural selection of mental functions has a name—evolutionary psychology: The aim of this discipline is to try to determine what factors might have driven the evolution, by natural selection, of our various mental capacities. Taking a clue from Kipling, critics like Stephen Jay Gould call the explanations proposed by evolutionary psychologists “Just So Stories”: they may sound good, but whether they are true or not can’t be known! After all, there is no fossil record of the mind, and evolutionary explanations about the mind are basically untestable. Gould suggests, in contrast, that language and perhaps other mental functions that are being touted as innate are not evolutionary adaptations, traits that were selected for, but rather are exaptations, features that enhance our fitness but

that weren’t originally designed to do what we use them for.” For example, birds probably developed feathers

as a means of controlling body temperature, and only later used them to fly. As Gould notes, “Once you build a complex machine, it can perform so many unanticipated tasks. Build a computer ‘for’ processing checks at the plant, and it can also..., whip anyone’s ass (or at least tie them perpetually) in tic-tac-toe. Another critic of the evolutionary explanation of language is the psychologist David Premack, a pioneer in the study of chimp language. As he notes, “Human language is an embarrassment for evolutionary theory because it is vastly more powerful than one can account for in terms of selective fitness. Evolutionary psychologists have responded to such attacks pointing out, among other things, that there are bad evolutionary explanations (those that are indeed “Just So Stories”) but also good ones (those that are better grounded). The problem is there’s a great deal of room (maybe too much) for interpretation in the decision about what constitutes a good as opposed to a bad evolutionary psychological explanation. A more general critique of evolutionary psychology can be found in *Alas, Poor Darwin*, a recent collection of essays put together by Hilary and Steven Rose (Steven is a neuroscientist who studies memory and a vocal opponent of nativist psychology). Regardless of how this ongoing debate ends, however, for now nativists are enjoying a good ride (though not a free one) in psychology.

Even if evolutionary psychology turns out to be dead wrong (and though I don’t think that’s the case, I am concerned about whether its explanations can be put to rigorous scientific tests), it wouldn’t mean that nativism is also completely wrong. It’s important to distinguish between the question “Is the function X innate?” and the question “Why did the innate function X evolve?” The former is a question about genetics, which can potentially be answered, while the latter concerns historical facts that are not easily verified scientifically. As Gould’s exaptation argument points out, something can be innate (passed on genetically) even if the role it plays today is not the one for which it evolved.

LEARNING DOMAINS.

Classic nativism, dating back to the Greeks, assumed the existence of a priori knowledge, of innate ideas, such as ideas about God, numbers, and other so-called universal truths. Modern nativists do not argue so much that we have a priori knowledge per se but rather that we have innate predispositions to acquire (learn about) specific kinds of environmental information.” For example, although nativists view language as an innate capacity in humans, they do not assume that we automatically know words and rules: to acquire a specific language we have to be exposed to people who speak that language. However, the information obtained from listening to others is, on its own, insufficient to fully account for language acquisition (especially the use of grammar)—it is this “poverty of the stimulus” argument that is key to the nativist assumption that language is innate. So Nativists further assume that the innate cognitive apparatus that makes language acquisition possible does not help us learn other things, which are managed by other independent innate mental modules. The capacities underlying innate mental modules have been described as domain-specific or information-specific learning biases, “specialized structures tuned to learn a particular kind of knowledge. The nativists’ modular view of learning is obviously inconsistent with the behaviorist notion of a universal learning capacity present in all species and used for all varieties of learning.

We now know, and will see in detail in later chapters, that the brain does indeed learn different things using different systems, and that different animals can have different learning capacities, which is consistent with the nativist view of innate learning modules and inconsistent with the behaviorist notion of a universal learning function. It seems possible, maybe even likely (though not yet proven), that the overall design of these systems is to a large degree programmed genetically. This would account for why members of a species share capacities, and why some capacities differ between species. At the same time, while the various learning systems may well have unique genetic histories, at the level of cells and molecules it appears, as we’ll see in the next two chapters, that relatively few solutions are used to achieve learning both within and across species. This implies that deep down there is a kind of universality to learning. We’ve encountered one of these mechanisms already, Hebbian plasticity, and will have more to say about it and related forms of plasticity in subsequent chapters.

Most systems in the brain can learn from experience, which means that the transmission properties of their synapses can be altered by experience. In this sense, learning itself is not the specific function that circuits were built (by evolution) to perform. Learning is instead a capacity of synapses, something that contributes to the way circuits work. The defense circuit involving the amygdala, for example, exists to detect and respond to danger. Learning is a very useful feature in a circuit like this, allowing it to respond to stimuli that are not in themselves dangerous but that have come to be associated with danger in the past—the sound made by a predator prior to an attack, for

ample. Without the capacity to learn from experience, the fear system would still be able to do its job (that is, to detect and respond to danger) but would be limited in that it could only respond to preprogrammed dangers.

SELECTIONIST NATIVISM VS. INSTRUCTIONAL CONSTRUCTIONISM

How, then, might an innate circuit, one used for learning a certain kind of knowledge, be wired during development? As we've seen, selection and instruction are the two main choices. Psychological or cognitive nativists reject on principle the idea that the environment "instructs" the mind, and therefore reject environmental instruction of circuits as well. As a result, they tend to be Selectionists. Massimo Piattelli-Palmarini, for example, builds upon Jerne, Changeux, and Edelman in denouncing the empiricist implications of instruction, and argues that "there is no such thing as learning." He is not questioning "learning" in the general sense of the term but in the behaviorist sense of being stamped in by the environment. He goes on: "There is no known process, either in biology or in cognition, that literally amounts to learning in the traditional 'instructive' sense, that is, to a transfer of information from the environment to the organism. . . . all the mechanisms of acquisition . . . are due to a process of internal selection." In a similar vein, Jacques Mehler has stated that knowledge comes from "unlearning," implying that we come to know things by a process of elimination.

Nativist psychologists assume that all human capacities are present in the exuberantly wired brain, waiting to be picked, or parameterized (to use Piattelli-Palmarini's term). However, as we've seen, synaptic selection alone can't fully explain brain development—instructed growth also plays an important role.

Nativists might at this point counter that the evidence for activity-induced growth is strongest for sensory systems, leaving open the possibility that the brain systems involved in higher cognitive processes (the mental modules of interest to nativists) might be built by selection. However, until it can be shown that higher cognitive functions are built by selection alone (that activity does not induce growth in areas involved in higher cognition), we have to assume that the basic principles discovered so far for sensory systems apply to other systems—that activity-induced growth is a viable possibility in cognitive systems.

Not all cognitive scientists are nativists—some, in fact, strongly oppose nativism. For example, in 1997, two influential publications attacked nativism head-on. One was *Rethinking Innateness*, a book by Jeffrey Elman and colleagues; the other was a journal article by Steven Quartz and Terry Sejnowski titled "The neural basis of cognitive development: a constructivist manifesto." Both publications argued for a *constructivist* approach to neural development, one that places instructed growth in the spotlight and emphasizes cortical plasticity rather than cortical specificity. While constructivists are especially opposed to the idea of innate knowledge (innate mental content, like words or concepts), they are generally displeased with most notions of innateness, including the innateness of capacities (like the capacity to learn a language).

A major tenet of constructivism is the idea that the wiring of cognition in the neocortex can be achieved on the basis of extrinsic inputs—that neocortical circuits can be constructed by extracting structure from the sensory environment. Constructivists accept the principle that genes build the overall structure, and that regressive events like synaptic selection take place, but they argue that epigenetic prewiring and selection alone cannot explain cognition. In addition to citing evidence for activity-induced growth of synapses, they refer to studies showing that damage to areas involved in language processing in the brain of a child leads to only minor deficit, if any, whereas the same lesion in

an adult causes severe problems in using language. Thus, in early life, before cortical circuits are finalized, areas normally not involved in language can take on language functions. Similarly, they cite the transplantation studies discussed earlier in which pieces of cortex are removed in young animals and put in a different location in another animal's brain. (Recall that the relocated bit of cortex then acquires the mental properties of the host.) Findings like these lead constructivists to conclude that circuits are not genetically programmed but acquire their functions from the neural activity that the environment elicits in them. At the same time, constructivists accept that there are constraints on how synapse construction works in the cortex. For example, cortical areas develop in conjunction with subcortical circuits, which are believed to be more subject to innate assembly than cortical ones.

The constructivist position accounts for a good deal of the data on cortical plasticity and is consistent with the notion that learning may be accomplished in fairly universal ways at the cellular level. Although I have considerable sympathy for much of the constructivist argument, I feel that it, too, has several limitations. First, it assumes that mental development can be understood more or less exclusively from the point of view of the cortex; noncortical contributions, though acknowledged, are viewed as mostly hardwired, and not given much consideration. However, because cortical areas depend on subcortical areas for the relay of environmental information, hardwiring

in a subcortical circuit will strongly influence what happens in a related cortical area, especially if the subcortical circuit matures earlier, as is often the case. The information ultimately extracted by the cortex is thus as much due to the organization of the subcortical circuit as to the structure of the environment. Second, in arguing against the innateness of some capacities, especially language, the constructivists may be throwing the baby out with the bathwater for the sake of consistency. Subcortical circuits are clearly more likely to be hardwired than cortical ones, but we shouldn't reject outright the possibility that some cortical circuits may also be hardwired. (For example, it has been proposed that the perception of facial expressions of emotion is performed by a species-specific face perception module in the primate, including human, cortex.) Third, in proposing that the cortex is a general-purpose cognitive learning device, not enough credit is given to the multiplicity of systems (both cortical and subcortical ones) that engage in learning. This throwback to universal learning just won't do. Fourth, in emphasizing the ability of a piece of cortex to adapt to new circumstances (due to deprivation or damage or surgical transplantation), constructivists overlook the fact that the adaptation that results from such unusual conditions may reveal abnormal rather than normal function: just because the color-processing part of the primate cortex can take on a new function when the color-processing area of the thalamus is damaged or removed does not contradict the possibility that color vision is an innate specialization of primates. Plasticity, again, is a feature of brain systems, not their evolved function.

A BOTTOM LINE. Clearly, more research rather than more speculation is in order at this point. However, armed with the information we've considered, let's return to the question posed earlier: Where does the epigenetic nature of synapse development leave innateness? It seems that there is no good reason to reject the possibility that the overall function of some circuits is more or less innately determined by genetic coding, especially if we are talking about innate predispositions to acquire certain kinds of information rather than about innate knowledge per se. At the same time, there seems to be no reason to assume that the epigenetic expression of an innate function has to involve *only* selection from preexisting synapses in finalizing the circuit: selection and instruction probably work together to shape the final connections of most circuits. The innateness of species-typical characteristics is thus not at all incompatible with the epigenetic nature of brain development (that is, the fact that genetic and nongenetic factors interact), so long as we do *not* insist that specific knowledge is innate or that nongenetic factors (like environmentally driven neural activity in neural systems) influence only the selection of preexisting synapses. Further, we have to remember that while innateness (genetic coding of an existing function) can mean adaptation (natural selection of that function), this is not always the case.

INNATENESS II: INDIVIDUALITY

One logical assumption might be that if a capacity is innate at the level of the species, individual differences in that capacity should also be innate. Suppose, for example, that after a thorough analysis of the problem, scientists were to come to the conclusion that humans have a specific capacity called intelligence, and that this capacity is innately coded in the human genome (by the way, neither of these propositions is generally accepted). Would we then be justified in assuming that differences in intelligence between individuals or between groups of people are due to genetic variation between individuals? Not necessarily!

The complex relation between the innateness of individual and species characteristics is nicely illustrated by a study performed on two groups of garter snakes living in California.⁹⁰ The snakes are members of the same species, but some live in the marshy coastal areas and others in the drier inland regions. The coastal snakes eat slugs; the inland ones don't. Steven Arnold was interested in whether this difference in food preference was due to genetic or environmental causes.⁹¹ He took eggs from snakes living in both regions, and when the babies hatched, he immediately put them in separate cages, isolated both from littermates and mother and from their natural habitat. A few days later, he offered slugs to snakes from the two groups. The coastal babies ate the slugs, while the inland ones ignored them. He then placed essence of slug on cotton swabs and measured how often the snakes flicked their tongues at the swabs. (Tongue-flicking is how snakes take in smells.) Isolated coastal snakes flicked their tongues vigorously, while the inland ones showed little interest. Even within each group, however, there was some variation from individual to individual. Arnold therefore examined the heritability of the within-group trait by comparing tongue-flicking scores of siblings to see if they were more like one another than unrelated snakes. If related individuals were more alike than unrelated ones, then genes would be implicated. It turned out that family relations had little to do with the way the snakes responded. Slug preference *within* the population is thus not due to genetic causes. Arnold then crossbred coastal and inland snakes and found

that this led to more variability in the behavior of the offspring than in the "pure" groups. Mixing the genes of the two populations reduced the genetic influence on the behavior, suggesting that the differences in the behavior *between* the groups is genetic in origin. Together, these findings lead to the conclusion that genes help in the explanation of the behavioral difference *between* groups but not *within* groups, at least in this case.

The snake study suggests that even if our capacity for some mental function like intelligence or language or emotion is innate, we are not necessarily able to explain differences between people in these capacities in terms of genes. Genes may play an essential role in placing a function in the brain of every human, and at the same time make a relatively small contribution to differences in the way that function is wired in individuals. The only way to know for certain is to do the proper experiments. Indeed, for some behaviors in some species, crossbreeding and heritability studies do show a significant relation in the behavior of siblings, including those reared separately, which suggests (but, as we'll see below, does not conclusively prove) a genetic contribution to behavior at the level of individuals and their families.

Clearly, the most compelling case for a role of genes in individual human behavior comes from studies of identical twins separated at birth. Because their common postnatal environmental influences are minimal, this situation is believed to be the best test for a role of genes in shaping personality traits of an individual. Some remarkable observations have been made. For example, one pair of identical twins separated shortly after birth ended up in very different situations—one grew up as a Catholic in Nazi Germany, and the other lived as a Jew on a Caribbean island. In spite of being raised under vastly different environmental conditions, the pair ended up with peculiar shared traits, including reading magazines back to front, and keeping rubber bands on their wrists.

Though interesting, it's hard to extrapolate from these idiosyncratic findings. However, large studies of many pairs of twins have been conducted to assess their similarity in personality traits and intelligence. This was done by computing heritability scores, which estimate how much of a given trait is due to environment and how much to genes, for the different measures. Such studies have shown that while identical twins, living together or apart, are more similar than fraternal twins or nontwin siblings, they are by no means indistinguishable. In fact, the correlation between

pairs on any particular measure was at most only about .50 for identical twins living together, and was a bit less for those living apart, suggesting that genes account for at most half of a given trait.

It has been argued that the contribution of genes is overestimated in twin studies in at least two ways. First, an assessment of the genetic component of heritability does not distinguish the direct effects of genes from indirect effects: a child's inherited timidity may lead his parents to reassure him constantly, his peers to pick on him, and his teachers to ignore him, each of which might in turn reinforce his timidity or bring about other behaviors.

Second, heritability scores do not accurately contrast genes with environment, for things that are similar in the environments in which separated twins live inevitably get mixed in with the genetic contribution. There are many opportunities for twins to have shared experiences of one kind or another. The renowned developmental psychologist Howard Gardner notes the following: "For at least nine crucial months, the twins share the same environment—the womb of the birth mother. ... Then, too, they may not have been separated right at birth. (And under what extraordinary circumstances does such separation occur?) They may or may not have been raised for a while by family members. The children are not randomly placed; in nearly all cases they are raised within the same culture and very often in the same community, with similar social settings. Also, infants who look the same and behave the same may elicit similar responses from adults.

In the earliest days of life, as we've seen, the main environmental factors that could affect the developing twins are chemicals passed from the mother's bloodstream to the fetus—hardly trivial influences. Also, many connections form in the brain prenatally, allowing the twin fetuses to begin to receive information from the external environment in utero. For example, by the thirtieth week of pregnancy, the heart rate of the fetus changes when sounds occur in the environment² The fetal brain can even discriminate environmental events and can learn and retain information about environmental stimuli. And after birth, as Gardner noted, they are often not separated immediately, and frequently end up in similar environments, where tendencies to look or act alike could help convert subtle into more powerful similarities. The case of the twins described above shows that this is not always true, but heritability studies live and die on group averages, not on interesting observations of individuals. When identical twins, separated at birth, end up sharing characteristics as adults, we need to wonder whether this is due to their common genetic heritage, to common influences within the womb, or to subtle environmental similarities that shape the development of their synaptic connections. At the same time, we also need to place the similarities in the context of the many ways in which twins differ from one another.

In such quantitative analyses of groups, the goal has been to determine how the variability of a trait relates to genetic similarity (identical vs. fraternal twins, fraternal twins vs. siblings, children vs. parents) in the general population, rather than in individuals. But in the last few years, research on genes has exploded, and new approaches are taking over. The older behavioral genetics approach had as its aim distinguishing the contribution of the environment from that of genes (that is, of all genes taken together). Complex mathematical analyses were used to extract subtle differences in the contributions of genetic and environmental factors to variation in a trait in large groups. This approach is giving way to efforts to identify *specific* genes that might be involved in specific traits. The entire genome has been mapped in some animals, including humans, which opens the door for some very sophisticated studies of the role of particular genes rather than of genes in general. For example, genetically altered mice can be created that have a single gene missing. Depending on what gene is knocked out, dramatic differences in behavior can result. Elimination of the gene for a certain enzyme used in regulating the effects of calcium or other molecules inside brain cells, for example, can interfere with Hebbian plasticity at synapses and thereby disrupt long-term memory, as we will see in chapter 6.

But as important as genes are, they are usually contributors to traits rather than dictators. And seldom is any brain function as complex as those involved in behavioral or mental processes under the control of only one gene. The fact that disruption of one gene can lead to a memory deficit does not mean that memory is due to that gene alone. All effects of genes are expressed epigenetically, by way of interactions in the internal chemical environment between the proteins made by multiple genes, or by way of external environmental stimulation that elicits synaptic

activity that then induces genes to make proteins. The proteins made by genes can in turn modulate neural activity at synapses. Assembly of the individual is a big job, and genes have lots of partners in the process.

TIMES FOR LEARNING?

Many psychological theories of development propose that the child matures in stages. Freud divided them into the anal, oral, phallic, and genital. Piaget talked about sensorimotor, preoperational, concrete operation, and formal operation stages, while other theorists have proposed different terms. Underlying the unfolding of any of these developmental stages in behavioral and mental processes must be neural changes. And, as discussed earlier, there's plenty of evidence that the brain does go through all sorts of changes during development—synapse numbers are rising and falling, cells are being born and dying, dendrites are growing, brain energy use (metabolism) increases. The problem is that it is not clear how the changes at the neural level relate to those at the psychological level.

A movement called brain-based education attempts to use information about brain development and function to guide education policy and practice. One of its basic principles is that there are special times for learning—critical periods windows of opportunity—times during which information has to be obtained or it won't be stored.

A good deal of evidence exists for critical or sensitive periods. As we've seen, if a kitten's eye is closed during a narrow time span in early life, it won't develop proper binocular vision because binocular cells in the cortex don't mature normally. The effect is reversible if the eye is opened within the window of opportunity, but not afterward. This is why misalignment of a child's eyes has to be corrected in early life. Once the critical period is over—that is, once synapses in the cortex are wired—the window closes. Another example is birdsong. If birds that sing aren't exposed to their species-typical song during the sensitive period (which begins in the weeks after hatching and continues until sexual maturity), they can't sing the song as adults. If they are exposed only immediately before the critical period (when their synapses are not ready) or afterward (when the synapses are already wired), they cannot learn to sing. This is somewhat similar to what happens in language development: language learning is more flexible in children than adults, as anyone who has tried to learn a second language after puberty knows.

But not all forms of learning are subject to critical periods. And because different systems of the brain are involved in learning different things, there's no way to generalize from a few instances to all forms of learning. Whether math or music learning has a critical period for synaptic development is not known. However, as far as education is concerned, these may not be the most relevant questions, at least not at this point.

John Bruer, a leading thinker in educational practice, has cautioned that brain-based education is an idea whose time has not yet come, in spite of the fact that it is a thriving approach. Bruer argues that we simply don't know enough about how changes in the brain (such as changes in synapse number or density, or changes in brain energy metabolism) relate to learning, or at least not enough on which to base something as important as education policy.

Instead, Bruer argues, educators should make the most of psychological findings in their efforts to improve the educational process. Once we know more about how specific capacities are used and when they appear, we might then gain clues about what to look for and where to look for it in the brain.

An excellent example of how psychology and brain science can interact in a way relevant to education comes from the work of Paula Tallal and Michael Merzenich. Tallal, a cognitive psychologist who studies language, discovered that dyslexic children are deficient at detecting and discriminating rapidly changing sounds, and that the ability of infants to perceive such stimuli predicts later language skills. When she played the same sounds to dyslexic children at a slower rate, they perceived the sounds much more accurately. Tallal's work caught the attention of Merzenich, a neuroscientist who had been studying the ability of the neocortex to adapt to environmental stimulation, showing, for example, that giving monkeys extensive experience with a particular sound increases the area of cortex that processes that sound (this is described in detail in the next chapter). Tallal and Merzenich collaborated to create special video games that would expose children with dyslexia to certain carefully selected stimuli in an effort to

improve their ability to process spoken and written language. By starting with slow stimulus-presentation rates and building up toward the rate at which speech and reading normally occur, they hoped to take advantage of the plastic nature of the cortex to actively shape its synaptic organization by training it to process stimuli more effectively. It is too early to tell how effective their program is, but regardless of its ultimate outcome, they are to be praised for their efforts to take basic research from the laboratory and use it to help children with reading disorders learn to read. Perhaps there are also lessons from their studies that might suggest how neuroscience research can also be used to improve the way children without learning problems are educated.

This work is at the heart of the debate between those who believe in specialized learning modules vs. more general learning capacities, since Tallal and Merzenich are using nonlinguistic stimuli to try to influence language learning. Those who believe that language is an innate module with its own learning rules argue that nonlanguage stimuli should have no effect. So not only will the outcome of this work have practical implications, it may also be very relevant to fundamental issues about brain organization.

John Bruer's most recent book attacks what he calls the myth of the first three years. He argues that too much emphasis is given to the idea that specific kinds of learning have to take place very early in life, or the brain will never again be able to learn that kind of information. Learning is indeed a lifelong process. And as Alison Gopnik and colleagues argued in another recent book, every time the infant learns something, his or her brain is changed in away that helps it learn something else. In a review of this book, a prominent developmental expert, Mark Johnson, noted that the early years are crucial not because the window of opportunity closes but because what is learned at this time becomes the foundation for subsequent learning. Indeed, much of the self is learned by making new memories out of old ones. Just as learning is the process of creating memories, the memories created are dependent on things we've learned before.

MOVING ON

Learning and development are two sides of the same coin. We can't learn until we have synapses. And as soon as synapses start forming on the basis of intrinsic commands, they are subject to being influenced by our worldly experiences. Genes, environment, selection, instruction, learning—these all contribute to the building of the brain and the shaping of the developing self by wiring synapses. Although the extensive plasticity that is present in early life eventually stops, our synapses do not stop changing, but remain subtly changeable by experience. It's time to see precisely how this happens.

CHAPTER FIVE

ADVENTURES IN TIME

YOU HAVE TO BEGIN TO LOSE YOUR MEMORY, IF ONLY IN BITS AND PIECES, TO REALIZE THAT
MEMORY IS WHAT MAKES OUR LIVES. LIFE
WITHOUT MEMORY IS NO LIFE AT ALL. . . OUR MEMORY IS OUR COHERENCE, OUR REASON, OUR

FEELING, EVEN OUR ACTION. WITHOUT IT, WE ARE NOTHING.

—Luis BuSuel

Memory is a marvelous device, a means of transporting ourselves to earlier times. We can go back a moment, or most of a life. But as we all know, it's not perfect, and is certainly not literal. It's a reconstruction of facts and experiences on the basis of the way they were stored, not as they actually occurred.' And it's a reconstruction by a brain that is different from the one that formed the memory. Sometimes, details are lost, but the gist is there. At other times, we just can't come up with what we're looking for, though we know we once had the information. Occasionally, we remember things that didn't actually happen. Fallible as it is, however, memory usually does a pretty good job for us. Adelle Davis's indispensable guide to sixties dietary habits told us, "You are what you eat." But Bunuel's lament is probably more accurate: we are our memories, and without them, we are nothing.

But what is memory? To most people, it is the ability to consciously recollect, to remember what happened days, weeks, or years ago. This is what psychologists call explicit or declarative memory'. The information it summons is explicitly available for conscious recollection, and can be verbally stated or deduced. This kind of memory is extremely flexible, allowing, for example, the sound of two cars crashing together to trigger a vision in your brain of an accident you once experienced. Through explicit memory you can recall a phone number, the way someone looks, what you had for lunch yesterday, or what you did on your last birthday. It is this memory that is so savagely attacked in Alzheimer's disease. Though extremely important in our lives, it is, in fact, only one kind of memory. My mother, an Alzheimer's victim, can't consciously recall much of anything that happens to her, but she can flawlessly execute the Cajun tune "J'ai Passé D'avant Ta Port" on her accordion. The kind of memory that allows her to do this is called implicit or nondeclarative memory. Such memories are reflected in the way we act more than in what we consciously know. Underlying this psychological difference is the fact that different brain systems are involved in implicit forms of memory, on the one hand, and conscious/explicit/declarative memory, on the other. To state this in the context of terms used in the last chapter, implicit memories are formed by systems that engage in domain-specific learning, whereas a domain-independent system forms explicit memories.

In chapter 2, I presented the outlines of a theory of the self that built on this view of memory. I proposed that the self is in part made and maintained by memory, and that both implicit and explicit forms are involved. In this chapter, we'll look at what is known about the neural circuits underlying explicit and implicit memory capacities. The synaptic mechanisms of memory are the topic of the next chapter.

IN SEARCH OF THE ENGRAM

In 1904, Richard Semon, a German scientist, coined the term *engram* to refer to the neural representation of a memory. Two decades later, Karl Lashley, an American psychologist, had begun a search for the engram, the locus of memory in the brain, a quest that occupied him for much of his career.

Lashley was one of the pioneers in the use of behavioral tasks, such as mazes, to investigate the relation between brain and behavior in animals. In his now-famous studies, he trained rats to "run a maze." Once the habit was learned, he then examined the effects of removal of different parts of the neocortex, which was believed even then to be where cognitive processes live. The animals were subsequently tested to see if they remembered the task.

While these lesions sometimes disrupted memory, the impairment was not systematically related to the location of the damage. Instead, it seemed that the memory loss was related more to the size of the lesion than to its placement in the cortex.

The results suggested two principles to Lashley. The first was mass action:

the effects of brain damage on memory are due more to the amount of tissue removed than to the actual area removed. The second was equipotentiality:

different areas contribute equally to memory, and when one is damaged, the others can take over. Lashley's bottom line: the brain doesn't have a specific system devoted to the formation and storage of memory. In his view, memories are stored in a widely distributed fashion in the cortex.

In retrospect, Lashley's approach was faulty. He used a behavioral task that could be solved in different ways, and made lesions more or less randomly. The reason that bigger lesions had bigger effects was not that all areas made equal contributions, but instead that different areas made unique contributions. As the lesions got bigger, more processes were disrupted, leaving the animal with fewer options for solving the problem. And in many instances the deficits were due less to the involvement of the cortical areas storing memory than to their role in basic sensory or motor processes necessary to perform the task. For example, a small lesion in the back of the cortex disrupted vision, but blind rats could solve the maze in other ways, using their senses of touch or smell. Bigger lesions left the animal with fewer sensory capacities, and the deficits got bigger.

Today, researchers go about things differently. They use knowledge about the organization of the brain to guide their attempts to find memory circuits. But hindsight is always visionary. The organization of the cortex was not well understood in Lashley's day, so he took the only approach available, one that now seems crude. Lashley didn't find the engram, but he made the search for it a legitimate enterprise.

THE SEA MONSTER SURFACES

In 1950, Lashley proposed somewhat facetiously that learning was not possible, given what he had discovered (or failed to) about the brain. Unfortunately, this was the state of things at the time. But soon thereafter the tide turned when a young man named Henry underwent brain surgery in Hartford. Henry was one of several patients in whom parts of the temporal lobe were removed in an effort to control epilepsy. The temporal lobe is one of the four major lobules that make up the cerebral cortex, and is often the place in the brain attacked by epilepsy. Sometimes, temporal lobe seizures are so severe and uncontrollable that the only recourse is to remove the epileptic site. In Henry's case, temporal lobe areas were removed on both sides of his brain. After the surgery, his epilepsy was better, but the improvement came at a high cost: Henry lost his memory.

Henry, known in the memory literature as H. M., is probably the most famous case in neurological history; and he has been the subject of countless published studies on memory. Much of the initial work was carried out by Brenda Milner and her colleagues in Montreal. They found that although H. M. could recall many of the events of his earlier life, especially those that occurred up to several years before the surgery, he was unable to form memories for experiences that occurred after the surgery. He could remember things for a few seconds (he had short-term memories), but he couldn't convert this information into long-term memories. He failed to recognize himself in a mirror, for example, but did recognize himself from old pictures. Though his IQ was in the normal range for his age, his view of who he was—his visual concept of his self—got stuck in the past, welded to his old surviving memories.'

The temporal lobe is a complicated structure with many subregions. An analysis of H. M.'s lesion, based on the surgical report, indicated that the main temporal lobe areas affected were the hippocampus, amygdala, and parts of the surrounding cortex.' By comparing H. M.'s lesion with those in other patients, it seemed that the hippocampus was the area damaged most consistently when memory deficits resulted.

The hippocampus, like many other brain areas, was named because its shape reminded early anatomists of something—in this case, a sea horse. The word comes from the Greek word *htpokampus*, which translates as "sea monster." In the early part of the century, the hippocampus was viewed as part of the rhinencephalon (which refers to

smell or olfaction). It later emerged as one of the major structures of the so-called limbic system, and came to be thought of as playing a crucial role in emotional functions.” Paul MacLean, the originator of the limbic system concept, conceived of the hippocampus as the seat of the Freudian id, a place in the brain where ideas could be mixed together and confused. MacLean located the id in the hippocampus because he thought that the primitive architecture of this structure could blend things in a surreal, dreamlike fashion, and would be ill-suited to participate in cognitive functions. But, by the end of the 1950s, studies of H. M. and other patients had led to the conclusion that the hippocampus was crucially involved in one of the most important cognitive functions of the brain—memory.

CANS AND CAN'TS

At first, it was thought that H. M. and similar patients had a global amnesia, a complete loss of memory, and that the hippocampus was indeed a general-purpose memory machine. However, it was soon noticed that amnesics could learn to perform well on some tasks that depended on prior learning.’ For example, Brenda Milner asked H. M. to copy a picture of a star viewed through a mirror. To do this, the movements had to be done in the direction opposite from the way they seemed they should be made. It’s hard at first, but eventually most people can do it. H. M. was no exception. He learned the task, and he retained the learning. But if asked about the drawing, he had no conscious memory of having made it. Suzanne Corkin of MIT then found that H. M. also improved with practice in another manual skill learning task—one in which he was required to keep a stick held in his hand on a dot spinning on a turntable. As with the mirror drawing task, the more times he did it, the better he got. His ability to form memories about how to make precise movements (called motor skills) seemed intact.

Neal Cohen later examined whether the spared skill-learning ability in amnesia was restricted to motor skills or might extend to what are called cognitive skills, the ability to get better at doing well-practiced mental tasks.” He showed that amnesics could learn and retain the ability to read mirror images of words (*egral* is the mirror image of *large*) just as they could learn to draw mirror images. He also showed that the patients could learn some complicated rule-based strategies required to play certain games or solve puzzles. As in other examples, although they could learn to use rules to play games, they later had no recollection of playing the games themselves.

Studies of a phenomenon now called priming also turned out to be very influential.” Elizabeth Warrington and Larry Weiskrantz demonstrated that amnesic patients were able to learn to recognize stimuli (pictures or words) on the basis of incomplete depictions of the stimuli. Each stimulus came in five versions, from very incomplete to complete. Initially, they had to see the more complete stimuli to successfully figure out what each was. But with repeated trials, they were able to recognize the stimuli on the basis of the weaker cues. This kind of learning, now called priming, might therefore be independent of the hippocampus. Later studies by Larry Squire showed that the instructions given to subjects made a big difference in whether priming occurs. For example, in one version of priming, the subject is given a list of words. Later, if asked to recall the items that were on the list, amnesic patients perform poorly. However, if instead of being asked to recall the items, they are given fragments of words (*mot*) and asked to complete them, like normal subjects, they are more likely to come up with words that were in the original study list (*motel*) than with other words (*mother*) that were not in the list. Although the amnesic patients had little conscious memory of the study list items, prior exposure to the list had clearly been retained at some level, since it affected their performance.

Weiskrantz and Warrington also showed that the classical conditioning of eye blink responses is preserved in amnesia. In this task, a tone is paired with an aversive stimulus (usually an air puff to the eye). After hundreds of trials, the tone elicits eyelid closure immediately before the onset of the air puff. This precisely timed response protects the delicate tissues of the eye from the air puff. Amnesic patients show normal eye blink conditioning, though they later have no memory of having seen the conditioning apparatus.

In the early 1980s, Larry Squire and Neal Cohen introduced the distinction between declarative vs.

procedural memory to account for memory functions of the hippocampus vs. other brain systems. Their proposal was that the hippocampus mediates conscious memory (memory that can be verbally declared) and that other brain systems mediate other forms of memory that are not dependent on conscious processes. They initially used the term *procedural memory* because at the time, many of the memory abilities that survived in amnesic patients involved the execution of skills that could be learned as rules or procedures. But as more examples emerged of things that hippocampal-damaged patients could learn and retain, it was clear that not all of these involved procedures. Procedural memory was therefore later renamed with the more neutral designation: *nondeclarative memory*. The terms *explicit* and *implicit* memory, suggested by Dan Schacter, essentially refer to the same memory processes as declarative and nondeclarative memory. Damage to my mother's hippocampus probably explains why she can't consciously remember the things that happen to her (her declarative or explicit memory system is gone), and lack of damage to the nondeclarative or implicit system that stores the skill of playing the accordion explains why she can remember how to play.

INS AND OUTS

The hippocampal contributions to memory, like the contributions of any brain area to any function, are made possible by the synaptic pathways. These deliver information to the hippocampus, process it within the structure, and transfer the outcome to target areas. Much is known about hippocampal anatomy, but we will only touch on some of the key points here.

As we've seen in previous chapters, information about the external world comes into the brain through sensory systems that relay signals to the neocortex, where sensory representations of objects and events are created. Outputs of each of the neocortical sensory systems then converge in the rhinal cortical areas, also known as the parahippocampal region, which integrates information from the different sensory modalities before shipping it to the hippocampus proper. . Studies by Wendy Suzuki, David Amaral, and others have elucidated many of the details of connectivity.

The hippocampus and parahippocampal (rhinal) areas constitute what has come to be called the medial temporal lobe memory system, the system involved in explicit or declarative memory. (However, they are not part of the neocortical temporal lobe and are not involved in its sensory-processing functions. They are instead examples of the so-called old cortex chap. 3]; this is why they are associated with the limbic system, which, as described earlier, was defined by the now-antiquated distinction between old and new cortex.)

Although within the hippocampus many complex circuits participate in the processing of incoming signals, there is a main-line circuit running through it called the trisynaptic circuit that is especially important. The trisynaptic circuit involves the relay of signals from the rhinal areas to an input area of the hippocampus (dentate gyrus) and from there to other areas (the CA₃ and CA₁ regions) and finally to the output region (the subiculum), which projects back to the rhinal cortex, closing the loop.

The connections between the hippocampus and neocortex are all more or less reciprocal. As a result, the pathways taking information from the neocortex to the rhinal areas and then into the hippocampus are mirrored by pathways coming out of the hippocampus to the rhinal areas and ending in the same neocortical areas that originated the inputs. In this way, cortical areas involved in processing a stimulus can, as we will soon see, also participate in the long-term storage of memories about that stimulus.

It's important to reflect for a moment on the nature of these connections and their implications for what the hippocampus and rhinal areas do. The rhinal areas serve as convergence zones, brain regions that integrate information across sensory modalities and create representations that are independent of the original modality through which the information was processed. As a result, sights, sounds, and smells can be put together in the form of a global memory of a situation. Without this capacity, memories

would be fragmented. Convergence zones also allow mental representations to go beyond perceptions and to become conceptions—they make possible abstract representations that are independent of the concrete stimulus.

While the primate neocortex has several cortical convergence zones, fewer such areas exist in the neocortex of other mammals. This may be an important clue to differences between the cognitive capacities of primates and other animals.

Because the hippocampus receives inputs from several convergence zones in the rhinal region, it can be thought of as a superconvergence zone. This no doubt accounts for why the hippocampus plays an essential role in our domain-independent memory capacity. It can form explicit memories about the implicit workings of many domain-specific systems, such as face- and language-processing systems, allowing us, for example, to form a memory that includes both what someone says and what he looks like.

NOMADIC MEMORIES

There are basically two kinds of amnesia: one is an inability to remember things that happened before the surgery or brain injury (called retrograde amnesia), and the other is an inability to form new memories (called anterograde amnesia). While we've focused so far on anterograde amnesia, both forms are present in patients with damage to the temporal lobe. Interestingly, the degree of retrograde amnesia is graded over time—memory is worse for events right before the surgery than for earlier ones. H. M. remembers his childhood, but not his present life.

The idea that amnesia affects recent more than remote memories is known as Ribot's law, after the French psychologist Théodule Ribot, who proposed that "the new perishes before the old." Larry Squire performed a seminal study of Ribot's law on patients who received electroconvulsive therapy (ECT) for depression, a procedure that often produces memory disturbances as a side effect." Squire tested patients before and after ECT treatment in 1974 on their memory of television shows. Before ECT, the patients had fairly reliable memory of shows from the early seventies and late sixties, with decreasing memory for earlier periods. This is quite normal, since we tend to remember recently learned things better than older ones. After ECT, though, their memory of shows from the early seventies was worse than before ECT, while memory for the earlier periods was unaffected.

The graded effect of retrograde amnesia is now believed to occur because the role of the hippocampus changes over time—the hippocampus is needed for memory storage initially, but its role decreases as time goes by. Why would the brain work this way? Why would memories be nomadic?

Jay McClelland, Randy O'Reilly, and Bruce McNaughton have proposed an answer, one inspired by results of a kind of computer simulation called connectionist modeling. These simulations have been very useful in attempts to understand learning. For example, it is known that when a connectionist model tries to learn about relations between stimuli, it does better when new information is incorporated gradually into the memory store rather than when it is put in all at once. So-called interleaved learning prevents new information from interfering with old memories. As a result, if the model is trained to recognize characteristics of animals like birds ("can fly") and fish ("can swim"), and then encounters the fact that a penguin is a bird that "can swim" but "can't fly," the result is very different when interleaved learning is used than if rapid learning is used. With rapid learning, the new information tends to result in a shift of the knowledge base such that both fish and birds come to be treated as swimming animals. But with interleaved learning, where the representation is built up slowly over many repetitions, the network gradually refines the representation of a penguin as a bird that swims but doesn't fly. In other words, the new information doesn't interfere with the knowledge base and is instead gradually added to it.

Many researchers believe that explicit memories are stored in the cortical systems that were involved in the initial processing of the stimulus, and that the hippocampus is needed to direct the storage process. For example, the creation of a memory for a visual scene involves the transfer of the perception from the visual cortex to the parahippocampal cortical areas and then into the hippocampal circuits. The processed signal, the memory, is then fed back through the parahippocampal areas back to the visual cortex.

According to the interleaved learning hypothesis, then, the memory is initially stored via synaptic changes that take place in the hippocampus. When some aspect of the stimulus situation recurs, the hippocampus participates in the reinstatement of the pattern of cortical activation that occurred during the original experience. Each

reinstatement changes cortical synapses a little. Because the reinstatements depend on the hippocampus, damage to the hippocampus affects recent memories, but not old ones that have already been consolidated in the cortex. Old memories are the result of accumulations of synaptic changes in the cortex as a result of multiple reinstatements of the memory. The slow rate of change of the cortex prevents the acquisition of new knowledge from interfering with old cortical memories. Eventually, the cortical representation comes to be self-sufficient. At that time, the memory becomes independent of the hippocampus.

Researchers like Jonathan Winson, Gyorgy Buzsaki, Bruce McNaughton, and Matt Wilson believe memory consolidation occurs during sleep,⁷ and specifically that it is during sleep that the slow interleaving of information into cortical networks takes place. Recent studies support this notion. For example, Wilson and McNaughton recorded the activity of neurons in the rat hippocampus. Using technically sophisticated procedures, they were able to identify precise patterns of cell activity in the hippocampus as rats explored a novel environment. Subsequently, when the rats went to sleep, the neural patterns seemed to be repeated in the hippocampus, as if the rats were dreaming about the places they had explored. This is an impressive finding. Although it has not yet been demonstrated that the hippocampal playback during sleep is actually read and used by the cortex, the existing data are consistent with the possibility.

Recently, the nomadic memory hypothesis has come under fire. Morris Moscovitch and Lynn Nadel have argued that the hippocampus always remains involved in memory storage.²⁷ However, with time, the memory trace comes to involve more and more brain regions, especially cortical regions, so that damage to any one area, like the hippocampus, fails to produce a deficit, since the other areas compensate. While this issue will no doubt continue to be debated in the coming years, a recent study has provided some important new evidence. Mice were trained in a spatial task that is usually impaired by hippocampal damage. Neural activity was then measured in the hippocampus at various times after training. Initially, neural activity was high in the hippocampus and was directly related to memory performance. (Mice with higher activity performed better.) With time, hippocampal activity decreased and came to be unrelated to memory performance, while cortical activity increased and came to be related to memory performance. Chalk one up for nomadic memory.

FACT AND EXPERIENCES: ARE THEY DIFFERENT?

So far, we have treated explicit or declarative memory as a single kind of memory capacity. However, in the early 1970s, the psychologist Endel Tulving proposed that long-term memories could be distinguished on the basis of what they were memories *about*.³⁰ He argued that episodic memories are about personal experiences (things that happened to you at a particular time and place), while semantic memories are about facts (things you come to know, but have not necessarily experienced). For example, you can know that deserts are hot and dry from personal experience, or by learning about deserts in school, without ever having been to one. Tulving suggested that episodic memory, which requires conscious recollection of the time and place of some personal experience, is particularly characteristic of humans, whereas semantic memory, being the simple storage of a fact rather than a personal experience, is within the capacity of many animals.

Recent studies by Faraneh Vargha-Khadem, Mortimer Mishkin, and colleagues support Tulving's distinction. Their research involved children in whom the hippocampus was damaged early in life. Remarkably, these children managed to attend mainstream schools and learn the basic facts that schools impart, in spite of the fact that they had poor memory for their own experiences. Taken at face value, the results suggest that the hippocampus is involved in remembering personal experiences, but not in remembering facts. This was unexpected by most researchers, since both semantic and episodic memory are lost in patients with temporal lobe lesions. However, most such patients have had damage to the hippocampus and the surrounding cortical areas (parahippocampal region), while the children studied by Vargha-Khadem are believed to have mainly sustained hippocampal damage.

Not everyone accepts the conclusion that the hippocampus is selectively involved in episodic aspects of declarative memory. Two prominent memory

researchers, Larry Squire and Stuart Zola, for example, have long promoted the view that the hippocampus and parahippocampal cortex (along with certain thalamic areas) constitute a single system that works as a declarative memory unit, making both semantic and episodic memory possible. They mention several problems with Vargha-Khadem's conclusions. One is that episodic memory, while poor, is still present in the children, and that even a little episodic memory capacity may be sufficient to support a good deal of semantic knowledge. As a result, the deficit may have involved both episodic and semantic memory, but the episodic deficit may be more readily detected. Further, they argue that Vargha-Khadem and colleagues are not justified in saying, on the basis of brain imaging, that the damage involved only the hippocampus. This can be determined only from an autopsy in which a detailed analysis of the brain is conducted, and not merely by taking brain scans from living persons. According to Squire and Zola, even if there was more of a deficit in episodic than semantic memory, the conclusion that one involves the hippocampus and the other doesn't would still be suspect.

Obviously, it is very difficult to pursue the fine-grained analysis of questions about the involvement of particular brain regions situated next to each other in studies of patients with brain lesions, for brain damage has little respect for anatomical boundaries of interest to scientists. In order to understand the detailed anatomy of memory, even of human memory, it is necessary to attack the problem through studies of experimental animals, the topic to which we now turn.

IN SEARCH OF H. M.

Soon after H. M.'s amnesia was discovered, researchers began trying to achieve what Lashley was unable to do—namely, to create amnesia by damaging the brain in animals. If this could be accomplished, it would greatly facilitate the study of memory from a biological point of view. With the hippocampus identified as the experimental target based on H. M.'s misfortune, it seemed that success would come quickly. However, efforts to produce amnesia in animals by damaging the hippocampus were not resoundingly successful at first. As late as 1970, well over a decade after Brenda Milner first described H. M.'s amnesia, Peter Milner (Brenda's ex-husband and a neuroscientist at McGill University in Montreal) wrote, "Unfortunately for the experimental study of the hippocampal amnesia effect, efforts to reproduce it in animals have so far proved unsuccessful."

A big part of the problem was that it was not understood initially that hippocampal lesions in humans affected only a certain kind of memory, explicit or declarative memory. In the animal studies, which were conducted under the influence of the behaviorist view that learning was learning (chap. 4), any task that measured the effects of prior learning was a legitimate way of assessing the general capacity for memory. We now know that only some kinds of memory depend on the hippocampus, and so only tasks that measure that kind of memory will be affected when the hippocampus is damaged.

In humans, it turned out to be fairly clear how to define hippocampal-dependent memory. If conscious retrieval was required, the hippocampus tended to be involved. But this criterion isn't really suitable for animal studies. We have no way of knowing whether animals have the kinds of conscious experiences we have, so from a scientific point of view, consciousness is not a very good notion on which to base a comparative approach to memory.

So how can we study hippocampal contributions to memory in animals? As Peter Milner's quote above implied, animal studies are important for working out the details and mechanisms. Over the years, two approaches emerged. One focused on studies of recognition memory in primates under the assumption that the best way to figure out why H. M. had the particular memory disorder he had was to find a hippocampal-dependent form of memory in a species close to humans. The other was the result of a serendipitous observation in a rat.

MONKEY BUSINESS

The history of primate studies of recognition memory is long and complicated. I reviewed it in some detail in *The Emotional Brain*, and will give only a brief summary of the story here. Most of the work has involved a test of

recognition memory called delayed nonmatch to sample developed by David Gaffan, an Oxford psychologist. In this task, a monkey sees one object and then, after a delay, is given two objects, one that was seen before and a new one. The correct choice, rewarded with a peanut or Froot Loop, is the new object (the one that does not match the sample). Many of the key findings using this task were obtained by Mort Mishkin and Betsy Murray at the National Institute of Mental Health and Larry Squire and Stuart Zola in San Diego, as well as Gaffan. Initial studies showed that damage to the hippocampus produced a significant deficit if the delay between the first and the second presentation of the objects was long enough. An animal model of H. M. seemed to exist. Then, it was shown that damage to the amygdala and

hippocampus together produced even more of a deficit. This model seemed even better, since H. M. had damage to both. It turned out, though, that the effects of the lesions were not due to damage to the amygdala or hippocampus, but to damage of the surrounding parahippocampal region, which was incidentally injured during the monkey brain surgery. This was still satisfactory from the point of view of modeling H. M., since he also had damage in these areas, but it raised problems for the view that the hippocampus is a key part of the memory system. However, studies of other patients who had damage restricted pretty much to the hippocampus (as determined by postmortem examination of their brains) showed that severe explicit memory deficits could result from hippocampal pathology alone. This suggested that the hippocampus is involved in human explicit memory, but that delayed nonmatch to sample is not a good test of the hippocampal contribution. It may instead be a better measure of the contribution of the parahippocampal cortex to recognition memory.

It is now believed by some researchers that the hippocampus and parahippocampal areas are separate components of a temporal lobe memory system). As we've already seen, sensory information comes into the hippocampus from the neocortex by way of the parahippocampal areas, and memories are established in the neocortex by way of the reverse connections. In sum it seems that the parahippocampal areas and hippocampus make unique contributions, and the delayed nonmatch to sample task reflects the parahippocampal more than the hippocampal contribution.

This work is relevant to the issue discussed earlier regarding the role of the temporal lobe in episodic vs. semantic memory. Perhaps delayed nonmatch to sample in the monkey is more like human semantic than episodic memory. If so, the distinction between the parahippocampal areas and hippocampus proposed by Vargha-Khadem would be supported. This would still leave open questions about what role the hippocampus itself plays. While more work will be needed to sort out exactly what the hippocampus does for memory, studies in rats have already begun to suggest some answers, as we now will see.

A PLACE FOR SPACE

The second approach to animal studies of explicit memory began more or less accidentally. In the early 1970S, John O'Keefe was a researcher at McGill University in Montreal. Brenda Milner, who did much of the initial work on H. M., was just up the street. So was Wilder Penfield, a neurosurgeon who found that conscious memories were elicited in patients when particular areas of the temporal lobe, like the hippocampus, were electrically stimulated in an effort to isolate the locus of epilepsy prior to surgery. Memory and the hippocampus were big topics in Montreal, and O'Keefe wanted to know more about how the latter made the former possible.

He had placed electrodes in the hippocampus of a rat in an effort to record the electrical activity in this structure during the neural coding of experience. Normally, projection cells (chap. 3) in the hippocampus fire fairly slowly, about once per second. But O'Keefe noticed that the firing rate in individual cells rose dramatically, up to hundreds of times per second, when the rat was in a certain location. As soon as the rat left that spot, the cell stopped firing. When it came back to the same spot, the cell started firing again. Because these cells seemed to be encoding the spatial location of the rat, O'Keefe called them place cells.

Sometimes, one finding can help make sense of others. Recall that in 1970 Peter Mimer pointed out that researchers had been unable to create amnesia in animals. He went on to note that some tasks were affected by

deficits more consistently than others. Why this occurred was not well understood at the time. However, one of the few tasks that hippocampal lesions seemed to impair with any consistency was maze learning, a task that is often solved by spatial cues. Maybe O'Keefe's place cells were responsible for memory performance in mazes. Indeed, in the mid-seventies, David Olton of Johns Hopkins University devised a maze-learning task (the eight-arm radial maze) that could only be solved using spatial information, and damage to the hippocampus interfered with performance of this memory task.

In the late seventies, O'Keefe, now in London, together with a former colleague from McGill, Lynn Nadel, wrote a famous book, *The Hippocampus as a Cognitive Map*, in which they proposed that the hippocampus is fundamentally a spatial cognition machine, and its contribution to other aspects of memory is secondary to its role in spatial processing. The book was a tour de force that established spatial cognition as a major area of study in neuroscience. The study of the hippocampus in spatial memory in particular became a thriving area of investigation.

Many different research groups have studied the role of the hippocampus in spatial processing. In 1973, Jim Ranck published a seminal study of hippocampal cell activity that has inspired much work since. Ranck, and especially his colleagues Bob Muller and John Kubie, have remained key players in the effort to understand place cell function. Also noteworthy is the research of Bruce McNaughton, Carol Barnes, Matt Wilson, and their colleagues. They have used technical ingenuity to push studies of place cells to the forefront of innovation. In one study, they recorded simultaneously from more than one hundred cells, showing how hippocampal neurons represent space and how the firing of these cells predicts where the rat will go next. As we saw above, they have demonstrated that the record of the rats' travels is encoded by place cells and played back during sleep, a finding that has been used to support the notion that the hippocampus slowly feeds new memories to the cortex during sleep. Another major contributor is Richard Morris, who developed a now widely used maze, the Morris water maze, in which rats use external cues to swim toward a submerged platform. Damage to the hippocampus disrupts this form of spatial learning. Together, the studies of hippocampal firing and the effects of hippocampal lesions on maze learning have unequivocally shown that the hippocampus is involved in spatial memory. Might it be a place for space?

RELATIONS

Although most of the evidence from rats for a hippocampal role in memory has involved spatial learning, many tasks that monkeys and humans with hippocampal damage fail on have nothing obvious to do with spatial processing. Does this mean that the rat hippocampus and the primate hippocampus are fundamentally different? Not necessarily. Damage to the hippocampus in primates does produce deficits in spatial memory, along with deficits in nonspatial memory. Spatial processing might be a specific example, one that is especially important to rats, of a more general capacity of the hippocampus.

Howard Eichenbaum has long promoted the idea that spatial memory is one specific example of what the hippocampus does. Its general function, according to Eichenbaum and Neal Cohen, is declarative memory. But this takes us back to the problem we encountered earlier. If declarative memory is conscious memory, how can we legitimately study it in animals? Eichenbaum and Cohen proposed that declarative memory should be defined not by its conscious aspects, but instead by the kind of processing requirements it makes. Specifically, declarative memories are relational. Thus, activation of a declarative memory leads to the activation of other related memories. As a result, declarative memories can be activated independent of the context in which they were established, and by stimuli other than those that were initially involved in the learning. Eichenbaum and Cohen propose that the anatomical architecture of the hippocampus allows it to engage in relational processing, and that the various kinds of tasks that implicate the hippocampus in memory all depend on relational processing.

In Eichenbaum's view, the hippocampus is involved in both the semantic and episodic aspects of declarative memory. In this sense, he sides with Squire and Zola in proposing that the declarative memory in general is the business of the hippocampus. However, while Squire and Zola propose that the hippocampus and parahippocampal

cortex form a unified declarative memory system, Eichenbaum suggests that only the hippocampus is involved in the flexible and relational processing that encompasses declarative memory, and that the parahippocampal areas are involved in the less flexible, nonrelational memory of individual component representations that go into declarative memory.

CONJUNCTIONS

The relational theory of hippocampal function is not universally accepted. The predominant challenge to it, other than the spatial theory, has been the view that the hippocampus binds stimuli together, blending the various elements of an experience (the way it looks, sounds, and feels) into a unified representation that is divorced from component elements. In this theory, the original elements, though, are lost, having been formed into conjunctions, while in the relational theory, the elements remain discrete and the hippocampus links them together. Consider an example. In the relational theory, your memory of a meal in a restaurant would involve the simultaneous binding of separate memories of the people you were with, the food you ate, and the general ambience of the restaurant, whereas the conjunction theory proposes that all of the elements are blended as a single memory of the situation.

Although the notion of conjunctions has been a part of many theories, a recent proposal by Randy O'Reilly and Jerry Rudy seems particularly promising. They built upon the earlier theory of McClelland, McNaughton, and O'Reilly, and on Rudy's earlier configural theory, which posited that individual objects are 'configured' in such a way to account for all of the objects in a single representation. The theory failed to hold up because such global representations were difficult to substantiate. O'Reilly and Rudy argue that the configural theory was indeed partly wrong because much research subsequently showed that animals with hippocampal lesions could learn conjunctive relations. In the new theory, though, O'Reilly and Rudy propose that the data can be explained by assuming that the spared conjunctive learning was performed by the neocortex, and that the neocortex and hippocampus learn using fundamentally different rules. The hippocampus learns conjunctions naturally and rapidly. When the hippocampus is damaged, the neocortex can be forced, through extensive training, to also learn conjunctions, but normally it does not do so. As McClelland, McNaughton, and O'Reilly had proposed, the neocortex, in contrast to the hippocampus, learns by slow, interleaved learning. Ongoing research is testing this new version of the configural theory in relation to spatial and relational theories.

Clearly, animal models have generated many interesting ideas about exactly what explicit or declarative memory is and how it is supported by the anatomy of the temporal lobe. Although we still don't have all the answers, impressive progress has been made. We can expect that progress in understanding the role of the hippocampus in memory processes will continue to be made in animal models, and that this information will help, in turn, to understand the nature of explicit memory in humans.

THE WAY YOU DO THE THINGS YOU DO

Though we have firsthand knowledge of our explicit memories (which are about things we were once aware of), many aspects of our outward behavior and inner life are controlled by brain systems that store and use information implicitly, that is, without our awareness of their operation. Implicit memories are reflected more in the things we do, and the way we do them, than in the things we know.

We've already encountered some examples of implicit memory. Almost by definition, these are memories that can be formed in H. M. and other patients who have amnesia due to damage to the hippocampus and related brain regions. For example, learned motoric and cognitive skills, priming, and classically conditioned responses can each be learned and performed in the absence of the hippocampus. They must be mediated by brain systems other than the one involved in explicit memory.

While explicit memory is mediated by a particular system (contained within circuits of the medial temporal lobe), many different systems in the brain engage in implicit learning. While these are often described as “implicit memory systems,” this is something of a misnomer. The systems that engage in implicit learning are not strictly speaking memory systems. They were designed to perform specific functions, like perceiving stimuli, controlling precise movements, maintaining balance, regulating circadian rhythm, detecting friend and foe, finding food, and so on; plasticity (the ability to change as a result of experience) is simply a feature of the neuronal infrastructure of these systems that facilitates their operation.

It's worth noting that John O'Keefe suggests the same basic scenario holds for the temporal lobe memory system. As we saw above, he believes that the hippocampus was built to process space, and that space processing demands a certain kind of plasticity that has been co-opted in the service of the more general capacity expressed as explicit memory in humans. In this sense, explicit memory is, in Stephen Jay Gould's terms, an exaptation (chap. 4).

Many of the neural systems that engage in implicit learning in humans have existed throughout much of the evolutionary history of mammals, and probably other vertebrates. They work unconsciously not because of some grand design to hide aspects of mental life from our sentient self but simply because their operation is not directly accessible by the conscious brain.

Systems that function implicitly contribute in important ways to our most characteristic traits. Each of us has his or her own style of walking, talking, and thinking. We hold our bodies in a certain manner when we are standing or sitting. We notice things that some others ignore, and ignore things that some people notice. The way we smile and the kinds of vocal inflections we use also help define who we are. The extent to which we are calm and collected, or emotionally reactive, when things go awry is also revealing, as are the logical paths and illogical leaps of thought we have. These and many other aspects of mind and behavior are expressed so automatically, so implicitly, that they may seem unchangeable, perhaps innate. But we should not overlook the crucial role of experience, which is to say of learning and memory, in establishing and maintaining them.

In recent years, considerable progress has been made in elucidating the neural circuits underlying certain examples of implicit learning and memory. While observations on H. M. and other humans led researchers to the hippocampus and its role in explicit memory, it has been studies of nonhuman species that have revealed the neural circuits of implicit memory. Studies in humans have, in many instances, verified that the same circuits are at work in our brains, but the basic findings have for the most part been discovered through studies of other animals.

THE RIGHT STUFF

Although the forms of implicit learning and memory that are now understood in most biological detail in mammals are each examples of classical conditioning, few studies before 1970 employed classical conditioning as a tool to explore the neural basis of learning and memory. Earlier research was inspired by the behaviorist tradition in psychology which emphasized operant or instrumental conditioning over classical conditioning. In classical conditioning, the subject learns an association between two stimuli—a bell and food, for example. The responses that result are automatic (e.g., salivation or heart rate changes). In contrast, in instrumental conditioning, the association is between a stimulus (reward) and response—food occurs when a rat presses a bar or turns the corner in a maze. The food reinforces the response, and, as a result, the subject repeats the response to get the food. The subject is passively involved in classical conditioning (the food comes when the stimulus occurs, regardless of what the subject is doing) but in operant conditioning, the subject is actively involved since the reward does not come unless the subject responds in a certain way. Instrumental conditioning was viewed as a more suitable means of accounting for the full complexity of human behavior, which the behaviorists clearly hoped to do. From Lashley onward, research on the brain mechanisms of learning and memory tended to emphasize the use of instrumental tasks.

By the 1970s, though, classical conditioning procedures had begun to take on a new life. Their rebirth began with the publication of an article by Eric Kandel and Alden Spencer in 1968. Noting the gap that existed between

learning (something that involves behavior) and plasticity (something that involves neurons and synapses), they proposed a step-by-step strategy that would allow the discovery of the neural basis of learning at the level of cells and their synaptic connections. This cellular-connection approach was the right stuff. It transformed the field, and is still in place today

Kandel and Spencer emphasized that the first step was to select an organism that expresses an easily measured and quantifiable behavior that changes with experience. Then, the neural circuit underlying unlearned and learned versions of the behavior should be identified. Next, the cells and synapses in the circuit that change with learning should be pinpointed. Finally the mechanisms that mediate the neural changes should be determined.

Kandel and Spencer's manifesto challenged researchers to think in terms of circuits rather than large areas of the brain. Lashley had searched for the locus of memory by making lesions in widespread areas, without any sense of where to search—he was really looking for a needle in a haystack. The cellular-connection approach told us how to pinpoint the part of the haystack where the needle might be hiding.

Kandel and Spencer believed that the cellular-connection approach could best be implemented in a simple nervous system containing a relatively small number of well-defined neurons. For this reason, Kandel chose to study behavioral conditioning and plasticity in snails. Since the snail doesn't have a hippocampus, the lack of involvement of the hippocampus in conditioning was not an impediment to the use of these procedures. The emphasis was on conditioning as a tool for studying behaviorally relevant neural changes rather than as a direct model of human memory. The payoff of this approach was fast. By the early 1970s, Kandel and colleagues had implicated specific neurons and synapses in learning in snails. We'll consider some aspects of Kandel's elegant and groundbreaking research in the next chapter, work that led to his being a Nobel Prize recipient in 2000.

The success achieved by Kandel and others in invertebrates using simple conditioning approaches to learning tempted researchers to pursue this approach in mammals and other vertebrates. The fact that the tools for studying neuroanatomy had just undergone a revolution greatly facilitated the effort to trace connections from start to finish (from sensory input to motor output systems) in a complex brain. Now, several decades later, much is known about the neural circuits underlying certain forms of classical conditioning. To the extent that these circuits do not depend on the explicit memory system, they are examples of circuits that engage in implicit learning.

But when this new wave of research on the neural basis of conditioning started in the 1970s, the distinction between explicit and implicit memory did not exist. Researchers pursuing the cellular-connection approach in invertebrates and vertebrates alike were using conditioning as a tool for studying the relation between neural plasticity and behavior. Only later was it realized that these studies were revealing the neural basis of what we now recognize as implicit memory

Today, we have a detailed understanding of circuitry involved in only a few of the many systems that learn implicitly. The examples discussed below are therefore not necessarily representative of implicit memory in general, but instead are representative of the implicit memory functions that have been characterized, in the tradition of the cellular-connection approach, from sensory input to motor output at the neural level.

BEING DEFENSIVE

One of the first researchers to take up the cellular-connection approach and apply it to the vertebrate brain was David Cohen. 'While recognizing the importance of studies of simple creatures, he felt that it was also important to pursue this approach in vertebrate species. His animal of choice was the pigeon, and his behavioral paradigm was Pavlovian conditioning.

Pavlov had been studying digestion in dogs when he noticed that they salivated when the person who fed them walked into the lab. He turned this into an experiment, and demonstrated that if a bell was rung just as food was put in the dog's mouth, the dog would salivate when it heard the sound of the bell, even if the food did not follow."

Cohen's research involved the same general procedure I've used in my work, Pavlovian defense or fear

conditioning. In Cohen's conditioning procedure, the appearance of a light was followed by a shock. As a result, the onset of the light eventually led to a change in the rate at which the pigeon's heart beat. The change in heart rate, as we've seen, is but one of many responses that occur during defense conditioning. Cohen simply used heart rate as a convenient way of assessing that defense conditioning had taken place. He made impressive progress in his cellular-connection studies, identifying input circuits through which the light was processed and output circuits through which the heartbeat was controlled. However, he did not succeed in connecting the input system with the output system.

It had long been known that the amygdala plays an important role in fear responses and even in learning to fear new stimuli." And one of the regions implicated in heart rate conditioning by Cohen was the bird-brain equivalent of the amygdala. However, Cohen's work fell short of demonstrating that the amygdala is the interface between inputs and outputs, and a likely site of the key neural changes that underlie fear learning.

Studies in the 1980s and 1990s painted a fairly comprehensive picture of the fear-learning circuits in the mammalian brain. Work from my own lab identified the way that inputs get to the amygdala, and studies by Bruce Kapp and his colleagues at the University of Vermont provided evidence for the output pathways exiting from the amygdala. Michael Davis, Michael Fanselow, and Norman Weinberger, and their students and colleagues, also made important contributions to various aspects of the fear-conditioning circuitry. Many of the details of the fear-conditioning pathways were described in *The Emotional Brain*, and some aspects of fear-conditioning circuits have already been discussed in earlier chapters. Here, I'll give only a brief summary of them.

The amygdala contains a dozen or so distinct divisions or areas, but relatively few are important for fear conditioning. The lateral nucleus of the amygdala (lateral amygdala) serves as the input zone. It receives information from the various senses, allowing the outside world to be monitored for threatening information. In studies performed together with Alsa Pitkänen, we found that the lateral nucleus has connections with most of the other amygdala regions. Karim Nader, Prin Amorapanth, and I then determined that only the connections with the central nucleus (central amygdala) are essential for fear conditioning. The central nucleus, the output zone, has connections with networks that control fear behavior and associated changes in body physiology. When the lateral nucleus detects some threatening stimulus, the central nucleus initiates the expression of defensive behaviors (like freezing) and other bodily responses associated with fear reactivity (changes in blood pressure and heart rate, stomach contractions, sweat gland activity, etc.).

Studies by Liz Romanski, Claudia Farb, Neot Doron and me show that the lateral amygdala gets inputs about the stimuli from two sources. It receives a crude but fast representation from a subcortical area (the sensory thalamus) and a slower but more complete representation from cortical sensory areas. The role of these two input systems to the amygdala was elucidated in studies performed in my lab by Liz Romanski. The path from the thalamus through the cortex to the amygdala, the so-called high road, allows complex information about objects and experiences to initiate fear reactions. But the amygdala also can be activated directly from the thalamus. Since this low road bypasses the neocortex, it only provides the amygdala with a crude representation of the external stimulus. But the arrival of crude information can have important consequences. For example, Fabio Bordi and I found that cells in the amygdala are able to determine the intensity or loudness of a sound through the thalamic pathway. Loudness is a good clue to how close something is, and distance is a good clue to how dangerous that thing is. If you treat loud things as dangerous, even if you don't know the source of the noise, you are probably going to be better off in the long run. So simply by computing intensity from the thalamus, the amygdala can immediately deduce significant details about a stimulus. Intensity is not the only feature gauged by the low road from the thalamus, but it's an important one.

The cortical route includes several more synaptic connections than the thalamic path to the lateral amygdala. Each synaptic link adds time to the transmission process, which is why cells in the lateral amygdala respond to information directly from the thalamus faster than they can respond to information from the cortex. More processing time by the brain means a slower mental and behavioral response from the organism. In situations where rapid responses are required, speed can be more important than accuracy.

Although inputs from the thalamic and cortical paths arrive on different schedules, the inputs reach the same

neurons. This means that in addition to jump-starting the system, allowing for rapid initial responses, the thalamic information can also prime lateral amygdala cells to receive the more exacting information from the cortex. As a result, the cells are then more capable of charging ahead if the cortical information confirms the threat, or of putting on the brakes if the cortical information establishes that no danger is present (for example, that the loud crackling sound was from you stepping on a branch rather than from something dangerous, like a bear about to pounce on you, or that the curved shape on the ground is a stick, not a snake).

Just because the cortical route allows a more refined stimulus analysis to trigger the amygdala, it should not be assumed that the high road is a conscious route to the amygdala. This was often misunderstood by readers of *The Emotional Brain*. The amygdala engages in implicit processing, including implicit learning, regardless of which pathway provides it with sensory information. As with any other stimulus, we become consciously aware of an emotional stimulus only when that stimulus is processed by networks involved in something called working memory, which we'll discuss in detail in chapter 7.

Using the cellular-connection strategy, a likely neural circuit required for fear conditioning has thus been identified from sensory inputs through motor outputs. Work is now under way toward achieving the goal of determining where the essential plasticity is in this circuit. While this research is still in progress, experiments have pointed to the lateral nucleus as a key site of plasticity in situations where neutral sounds are paired with shocks. First, studies by Liz Romanski found that cells in the lateral nucleus receive convergent inputs from pathways that process both sounds and shocks." Such convergence is believed to be essential for conditioning to occur. Indeed, Greg Quirk, Chris Repa, and Michael Rogan found that the activity of lateral amygdala cells increases when the tone and shock occur at the same time but not when they are separated. These changes in the neural responses precede the development of conditioned fear behavior, suggesting that the neural changes might well be responsible for the behavioral learning. Studies by Steve Maren and Denis Pare have now also shown plasticity in the lateral amygdala." Further, conditioning does not occur if the lateral nucleus is damaged. Conditioning is also prevented if synaptic activity and/or plasticity is disrupted by injecting certain drugs in the lateral nucleus, as shown by Jeff Muller, Ann Wilensky, and Glenn Schafe in my lab, as well as by research from other labs.⁷⁰ Together, these findings indicate that the lateral amygdala is a key site of plasticity during fear learning. As we will see in the next chapter, progress has also been made in uncovering some of the precise cellular and molecular mechanisms in the lateral amygdala that underlie fear conditioning.

As Pavlov suspected, defense conditioning plays an important role in the everyday life of people and other animals. It occurs quickly (one pairing of the neutral and aversive stimulus is often sufficient) and endures (possibly for a lifetime). These features have no doubt become a part of the brain's circuitry due to the fact that an animal usually does not have the opportunity to learn about predators over the course of many experiences. If an animal is lucky enough to survive one dangerous encounter, its brain should store as much about the experience as possible, and this learning should not decay over time, since a predator will always be a predator. In modern life, we sometimes suffer from the exquisite operation of this system, since it is difficult to get rid of this kind of conditioning once it is no longer applicable to our lives, and we sometimes become conditioned to fear things that are in fact harmless. Evolution's wisdom sometimes comes at a cost. Thompson and his students. Thompson had worked with Alden Spencer, and his studies of plasticity in the spinal cord had contributed significantly to the genesis of the cellular-connection approach subsequently developed by Kandel and Spencer. In the 1970s, Thompson began to study brain mechanisms underlying the Pavlovian conditioning of eye blink responses in rabbits as a simple mammalian preparation for pursuing the locus of learning in the brain. We encountered this kind of conditioning in our discussion of H. M.

Earlier studies by Isadore Gormezano and others had worked out in precise detail many of the basic principles underlying eye blink conditioning to a sound that was followed by an annoying stimulus (a shock or a puff of air to the eye). Before conditioning, the shock or puff causes the eye to blink but the sound does not. After conditioning, though, the sound elicits a blink, and the blink is precisely timed from the onset of the sound—the eye

closes right before the time when the puff or shock would occur. The blink therefore appears to be an adaptive response that protects the eye from irritation.

Building on this behavioral work, Thompson and other researchers pursued the neural basis of eye blink conditioning. It was discovered that removal of the entire neocortex had no effect on the ability of rabbits to be conditioned. In fact, eye blink conditioning could occur following removal of most of the forebrain, including the amygdala, hippocampus, and other areas. This led to the conclusion that the essential plasticity underlying this form of conditioning must be in the brain stem rather than the forebrain. Indeed, studies by John Moore, a former student of Gormezano's, found that damage in the lower brain stem did prevent conditioning, while at the same time leaving the response elicited by the unconditioned stimulus unaffected). The rabbits could therefore blink their eye in response to a direct stimulus normally, but could not learn to blink to the warning sound.

Through a series of studies involving the use of many different techniques, Thompson and others (including Joseph Steinmetz, Mike Maur, and Chris Yeo) have elaborated the brain stem mechanisms of eye blink conditioning. They have been led to the conclusion that the critical site in the brain stem is the cerebellum, a wrinkled mass of tissue that sits on top of the brain stem and that had long been believed to be involved in the control of posture and movement and that had previously been hypothesized to be a site of learning. The researchers showed that lesions of certain areas of the cerebellum prevented conditioning, and that the activity of neurons in these areas changed during the course of conditioning. Pathway-tracing studies determined that the cerebellum has neural connections that allow it to integrate the sound and the air puff, and to control the eye blink response. In fact, the eye blink can be conditioned by substituting the tone and the puff with direct electrical stimulation of the neural pathways that transmit the tone and the puff into the cerebellum. Although not everyone favors the view that plasticity is an important feature of cerebellar function, studies by Steve Lisberger using a different training procedure have also implicated cerebellar circuits in motor learning.

Why should we care about how rabbits blink? As with many approaches in neuroscience, it is important to judge this work in a broader context. The research is not just showing how the rabbit brain controls the rabbit eye, but instead really concerns how movements come under the control of extrinsic signals. Control of precise movements at just the right time is a fundamental aspect of many kinds of acquired behavioral skill, from learning to ride a bike, to hitting and catching baseballs, to driving cars, and many other things that we do every day. The cerebellum, as we have seen, is also involved in posture and balance, and both of these functions are modified by experience—astronauts, for example, readily adapt to the changed balance and postural requirements imposed by altered gravity. The way we hold our body at rest, the way we move, and the grace (or lack thereof) with which we do things contribute significantly to the physical image we project. They are part of who we are. Studies of the role of the cerebellum in motor learning are thus revealing clues about an important class of implicit learning capacities.

THE TASTE OF THINGS TO COME

Most people like sweet and dislike bitter-tasting foods. This preference is present very early in life—an infant will spit out bitter foods on the first encounter. Infants born with massive parts of the forebrain missing and rats with lesions of all of the forebrain and midbrain still exhibit taste preferences, suggesting that taste is mediated by the hindbrain. The early appearance of taste preferences suggests that they are built into the hindbrain during prenatal development and are possibly innate. Nevertheless, taste preferences are subject to modification by experience. Many adults enjoy eating broccoli and spinach, even if they were disliked in childhood. And some foods, once enjoyed, can become highly aversive if they make you sick.

Studies of food poisoning in animals, especially rats, have helped us understand a fundamental kind of learning, so-called conditioned taste aversion (CTA), a special form of classical conditioning. Farmers had long known that cattle avoid foods that make them sick, but this avoidance was not studied seriously until the 1960s, when John Garcia and his students found that rats made ill by injection of a mild poison strongly avoided the last food eaten be-

fore getting sick. What was particularly striking about this experiment was the fact that the avoidance would develop even if sickness was induced hours after the food was consumed. The finding violated one of the cardinal principles of learning—that associations form between stimuli that occur at the same time, and not between stimuli separated by hours. The fact that CTA violated this law was one of the major challenges to the universal nature of learning promoted by the behaviorists (chap. 4). And because these data were viewed by strict behaviorists as impossible, Garcia had a great deal of trouble publishing his studies. But he persisted and eventually succeeded.

CTA has become a major area of research in behavioral neuroscience. Although the circuits involved in it are not as well characterized as those that govern fear and eye blink conditioning, a good deal of progress has been made. Most of the work has used saccharin, a novel sweet taste, as the conditioned stimulus (this is like the tone in a fear-conditioning study). Rats prefer drinking saccharin-flavored water over plain water, but if they get sick after drinking saccharin, they will subsequently avoid it. Often the sickness is induced by an injection of lithium chloride, a nausea-producing substance that serves as the unconditioned stimulus (like the shock in fear conditioning).

Like other forms of classical conditioning, CTA requires that the neutral taste and the nausea-producing stimulus come together synaptically at individual neurons. The taste pathway from the tongue goes to a region of the hindbrain called the nucleus of the solitary tract, which is involved in taste preferences. The same general region of the hindbrain receives fibers from the gut, telling the brain about nausea and other gastrointestinal conditions. However, because different regions of the hindbrain receive the nausea and taste signals, conditioning is not likely mediated by synaptic integration in this area. The two parts of the nucleus of the solitary tract do project onto a common area in the midbrain called the parabrachial nucleus. Damage to this region, where taste- and nausea-stimulus processing overlap, prevents CTA from occurring. Furthermore, in normal animals, cellular activity elicited in this region by taste stimuli increases following exposure to the nausea-inducing stimulus, indicating that cells in the region are involved in the conditioning of taste by nausea. The parabrachial nucleus thus is likely to be an important area of plasticity underlying CTA.

But this is not the whole story. Removal of the forebrain, which leaves only the brain stem to run the show, disrupts the acquisition of CTA, in spite of the fact that the same procedure leaves taste preferences undisturbed. The parabrachial nucleus thus does not work alone in the mediation of CTA and instead depends on the forebrain areas. Anatomical tracing studies give clues as to which areas of the forebrain are important here. The parabrachial nucleus sends fibers to the taste area of the thalamus, which in turn sends fibers to the taste area of the cortex. Further, the parabrachial nucleus, taste area of the thalamus, and taste cortex all send fibers to the central nucleus of the amygdala. Using experimental lesions, it has been shown that damage to two of these regions, the cortical taste area and the central amygdala, interferes with CTA. The taste cortex is believed to be involved in detecting and dis

criminating novel tastes, which is an important aspect of learning about the consequences of new tastes. The role of the central amygdala is controversial. Some believe it is not required for learning, whereas others believe it is essential. Differences in the way CTA is established or tested may account for the different views. While more work is needed to resolve these issues, CTA has proven to be a useful model system for studying memory circuits, and has also yielded some important findings about the molecular basis of memory, as we will see in the next chapter.

WHENCE SUCCESS?

Early research on learning mechanisms was inspired by the behaviorist tradition in psychology. By the 1970s, though, researchers came to realize that less is more when it came to the choice of behavioral tasks used to study learning, and began taking their inspiration more from neurobiology, especially from the cellular-connectionist perspective. The emphasis therefore shifted away from complex instrumental learning procedures to simpler conditioning tasks, since, as Kandel and Spencer noted, it would be easier to discover the circuits and specific synaptic changes involved in simple than in more complex forms of learning.

The advantages of classical over instrumental conditioning as a starting point for a neural analysis of

learning are readily illustrated by the history of research on fear learning. Early work in this area emphasized avoidance-conditioning procedures, instrumental tasks in which the subject learns to avoid an aversive event, like foot shock, on the basis of a warning signal. It turns out, though, that avoidance conditioning begins as classical conditioning: the warning signal produces the avoidance response because it was first associated with shock. For example, a rat can be taught that if it runs from one end of a maze to the other when the tone comes on, the shock will be prevented. But the first step is that the rat has to learn, through classical conditioning, that the tone predicts the shock. Only then can it use the tone-shock association to learn the instrumental response of avoiding the shock by running through the maze in the presence of the tone. It makes good sense, in retrospect, that the way to figure out how the brain mediates avoidance conditioning would be to first understand how the tone-shock association is formed, and then to ask how the tone-shock association is used to establish avoidance. But this was not realized in the early days, and progress in understanding fear learning was slow until researchers decided to simply focus on the classical conditioning of fear.

Obviously, memory researchers will not be satisfied with explaining only the simple forms of learning. But armed with knowledge of how those simple forms work, we may well be able to take on more complex ones. For example, now that the classical conditioning of fear is well understood, it might be possible to figure out the neural basis of avoidance. This is an important avenue of research since the hallmark of anxiety disorders is pathological avoidance of anxiety-producing situations, which greatly constricts the lives of anxious people. As we will see in chapter 9, recent studies have made considerable progress in using knowledge obtained about the neural basis of fear conditioning as a stepping-stone to understanding avoidance learning.

CONSCIOUS AND UNCONSCIOUS MEMORY RECONSIDERED

What exactly do we mean when we say that the hippocampus is involved in conscious memory and that other forms of memory, which are independent of the hippocampus, are not directly accessible to consciousness? Does it imply that the hippocampus makes consciousness possible? Probably not, since hippocampal damage does not disrupt normal conscious awareness.

Examples of memory deficits produced by hippocampal damage in humans have indeed mostly involved failures in the ability to consciously retrieve information about past experiences. When tested on learning tasks that do not require conscious retrieval, like priming or conditioning, hippocampal-damaged patients basically perform acceptably. But when required to retrieve information consciously, often within the same task they performed well on implicitly, they do poorly. For example, as we've seen, though hippocampal-damaged patients respond to a tone previously paired with an air puff to their eye, they have no conscious memory of having been conditioned.

The importance of conscious retrieval in defining the hippocampal deficit was highlighted in a study performed by Larry Squire and colleagues. They had normal subjects and persons with amnesia due to hippocampal damage watch a silent movie. The subjects were asked to pay attention to the movie while ignoring sounds and air puffs that would occur throughout. This type of instruction typically biases subjects to form explicit memories about the attended event, and was aimed at keeping the tones and puffs in the background of awareness. The tones and puffs were delivered in two different ways. Some of the subjects received a standard procedure in which the puff occurred at the end of the tone. As expected from prior work, both the nor-

mal controls and the amnesic patients conditioned—they learned to blink in response to the sound—in spite of the fact that they were instructed to ignore the tones and puffs. For other subjects, the puff came some time after the tone had ended. This is called trace conditioning and is known from studies of rats and rabbits to be disrupted by hippocampal damage. Indeed, the patients with hippocampal damage did not condition in the trace task. What is particularly significant, though, is what happened to the normal controls. Some of them conditioned and some didn't, and whether they conditioned was closely related to whether they reported, after the experiment, having noticed a

relation between the tone and the puff. In other words, if they were conscious of the relation, they underwent trace conditioning; otherwise they did not. This suggested that trace conditioning requires awareness, and that the hippocampus makes this awareness possible.

One problem with this conclusion is that it strains the interpretation of the rat and rabbit data, which reveal the same effect: deficits in trace conditioning following hippocampal damage. If the human hippocampus is involved in trace conditioning because of its role in awareness, then the rat and rabbit hippocampus must also be involved in conscious awareness, since they, too, are required for trace conditioning. If not, the hippocampus is not really involved in awareness.

A recent study by Marvin Chun and Liz Phelps suggests the latter is the case. They tested normal subjects and amnesics on a task that required the identification of a single letter *T* in a sea of *L*'s. The *L*'s thus formed the background or context in which the target *T* appeared. Unbeknownst to the subjects, the context stimuli came in two varieties. Sometimes the *T* was amongst an arrangement of *L*'s that had come before, and sometimes the *T* was in a novel arrangement of *L*'s. Both normal controls and amnesics got better as the task wore on. However, the normals but not the amnesics benefited from the repetition of the background pattern: they were better at finding the *T* in the repeated pattern than in novel ones. Nevertheless, the subjects were not aware, when asked later, of the fact that the background repeated. The hippocampus, in other words, was required to process and learn about the unattended background stimuli.

Chun and Phelps were motivated to perform their study by the fact that studies of rats had long suggested that the hippocampus was involved in contextual processing. For example, as we'll see in chapter 8, when you condition a rat by pairing a tone with a shock, the rat becomes afraid not only of the tone but also of the box in which the conditioning takes place. The hippocampus, it turns out, is necessary for this kind of conditioning because it is able to put together many things at once, to relate or configure them into a context. This is not unlike the way O'Keefe or Eichenbaum or Rudy views the rodent hippocampus: it has a certain arrangement of neurons that allows the creation of a memory involving a complex set of stimuli. There is no need to call upon consciousness in the interpretation of the rodent work.

The Chun and Phelps finding strongly suggests that the human and rodent hippocampus work the same way, forming memories about the relations between stimuli. Regardless of the nature of the relations (whether they are unified configurations/conjunctions of stimuli, relations among different memories, or memories about spatial arrangements), the key function of the hippocampus is probably processing relations. If so, then the hippocampus might not be fundamentally different from other systems that store information. That is, it probably goes about its business implicitly. What is different about the hippocampal and other memory systems is instead the fact that the hippocampus is synaptically connected in such a way that its activity is available to the brain systems that mediate conscious awareness (which we'll consider later in chapter 7), whereas implicit systems are not connected in this fashion.

THE GOOD, THE BAD, AND THE UGLY

Memory is amazing. A simple thought can take you to a past time in your life. And often without thinking at all we "remember" to do what we need to do each day. But with the good come the bad, and even the ugly, aspects of memory, its failures. No one knows more about this than Dan Schacter, who put the terms explicit and implicit memory on the map. His recent book, *The Seven Sins of Memory*, lays out in graphic detail the ways memory fails us. Schacter's sins are transience, absent-mindedness, blocking, misattribution, suggestibility, bias, and persistence)

Transience is simply the inability to hold onto information. Absent-mindedness is our annoying capacity to fail to pay attention to what we are doing, as when you put your keys down while doing something else, and then can't find them because when you put them down you were mentally involved in some other activity. Blocking is the failure

to pull out that fact or name that is on the tip of your tongue. Then there's misattribution, as typified by the belief that a memory formed in one situation when it actually occurred in another. This is particularly important in the context of eyewitness testimony.

Suggestibility is also relevant to eyewitness testimony and to the topic of false memories implanted in the therapeutic process. Bias creeps into memory in several ways, one of which is consistency bias, which leads us to revise our memory of a situation to make it fit what we feel or think now. Finally, there's persistence, which on the surface sounds good. Emotional arousal makes any memory stronger, which can be good, but when the memory is of a traumatic experience, persistence can be debilitating.

Schacter emphasizes that the sins are less design flaws than features. That is, he claims that the sins are by-products of virtues—persistence, as we've seen, can be good. Regardless of the adaptive value, if any, of the sins, at least with knowledge of their existence we can try to protect ourselves from their negative consequences.

REMEMBERING WHO WE ARE

This chapter opened with a quote from the Spanish surrealist filmmaker Luis Buñuel, who said, "Life without memory is no life at all. . . . Our memory is our coherence, our reason, our feeling, even our action. Without it, we are nothing." Memory does indeed make us who we are. But it is important to remember that memory is more than just what we can consciously recall. It therefore follows that a loss of explicit memory, due to damage to the hippocampus, while devastating in many ways, would not eliminate personality. In Alzheimer's disease, for example, the hippocampus and related areas are the first to be destroyed. Even in the face of severe memory problems, its sufferers initially remain much the same person they always were—they walk and talk the same, and have the same basic habits and traits. As the disease spreads widely to other brain areas, those that function implicitly, personality begins to break down.

In order to be yourself, you have to remember who you are. Keep in mind though that the memories involved are distributed across many brain systems, and are not always or even mostly available to you consciously.

CHAPTER SIX

SMALL CHANGE

NOTHING ENDURES BUT CHANGE.

--He rac I tutS

A LIVING THING IS DISTINGUISHED FROM A DEAD THING BY THE MULTIPLICITY OF THE CHANGES AT ANY MOMENT TAKING PLACE IN IT.

Life is change, and the brain is a device for recording changes—for forming memories through learning. Learning and memory, as we've seen, fill in the details of who we are as we become a unique person. But what is the nature of the neural changes that constitute learning and memory? Most neuroscientists today believe that alterations in synaptic connectivity underlie learning, and that memory is the stabilization and maintenance of these changes over time. How, then, does experience actually change synapses, and what makes the changes last?

HEBB'S MAGIC

In an 1894 lecture to the Royal Society of London, Santiago Ramón y Cajal, whom we met in connection with the neuron doctrine, proposed that “the ability of neurons to grow in an adult and their power to create new connections can explain learning.” While this statement is often cited as the origin of the synaptic theory of memory, it was anticipated by a number of other ideas. For example, in the mid-1700s, the philosopher David Hartley suggested that mental associations (memories about the relation between stimuli) are the result of vibrations between nerves.^z More than a century later, William James, the father of American psychology, wrote in his famous 1890 textbook: “When two elementary brain processes have been active together or in immediate succession, one of them, on reoccurring, tends to propagate its excitement into the other. And in his medical research days, Sigmund Freud argued, “Memory is represented by the facilitations existing between...

rons. Still, these suggestions, including Cajal's, were only partially formed when compared to the view proposed by the Canadian psychologist Donald Hebb, in his 1949 book, *The Organization of Behavior*.

Hebb made many seminal contributions to scientific psychology including pioneering studies on perception, instinctual and emotional behavior, and intelligence. However, he is best known for his synaptic theory of memory, his fire and wire theory, which we encountered in chapter 4 in the context of brain development. And although this theory is revered by brain scientists today, Hebb apparently did not believe it to be particularly important, and certainly didn't consider it to be his best idea or his most significant contribution.

Hebb's notion, as you'll recall, is that “when an axon of cell A is near enough to excite cell B or repeatedly and consistently takes part in firing it, some growth process or metabolic changes take place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.” Let's expand this idea a little so we can see how it might apply to memory, and especially to a memory of the fact that two stimuli once occurred together.

In order for two stimuli to be bound together in the mind, to become associated, the neural representations of the two events have to meet up in the brain. This means that there has to be some neuron (or a set of neurons) that receives information about both stimuli. Then, and only then, can the stimuli be linked together and an association be formed between them.

We'll consider a simple example to see how Hebb's theory makes associations possible. Imagine a neuron A that is postsynaptic to two other neurons. One of these (S) is strongly connected and the other (W) is weakly connected to A. As a result, when S fires, A also often fires, whereas A is less likely to fire in response to activity in W. Further, imagine that each of the two neurons connected to A is involved in processing distinct stimuli. Because the connection from S to A is stronger than the connection from W to A, A is more likely to fire in response to external

stimuli that activate S than those that activate W. If on some occasion stimuli processed by both of the cells occur simultaneously, the weak input to A will likely be occurring at the same time that A is being fired by the strong input. As a result, according to Hebb's rule, the weak connection will be strengthened. A is now a place in the brain where the strong and weak pathways are related, such that activation of the weak one now has the same effect that only activation of the strong one had before. This sort of thing happens all the time in daily life. If you are walking on the sidewalk in front of your neighbor's house (weak stimulus) and his dog bites you (strong stimulus), you will associate the sidewalk where you were bitten with the dog and be less inclined to walk that way.

Today, neuroscientists use the term *Hebbian learning* to describe change in the connection strength between two neurons caused by the fact that the postsynaptic cell was active when presynaptic inputs arrived. Hebb realized that his idea was not completely novel, noting that others had proposed synaptic theories of memory, and some historians have suggested that Hebb be given less credit in light of past views. The fact remains, however, that Hebb developed a theory of synaptic strengthening to account for learning whereas most of the others simply proposed that such changes might occur. And, as we'll see later in the chapter, Hebbian learning actually explains how certain kinds of synaptic changes take place in the brain, and may in fact be a major way that memories are made.

SEARCHING FOR SYNAPSES

Around mid-century, synaptic change was a topic that was in the air. Hebb's book had been published, as had one by the great Polish neuroscientist Jerz Konorski. Konorski used the term *plasticity* to describe the ability of a neuron to be altered by experience and had proposed a theory of synaptic plasticity not too different from Hebb's. And by the early 1950s, a number of studies had successfully shown that repeated delivery of a brief electrical stimulus to a nerve pathway could alter synaptic transmission in that pathway—could, in other words, produce synaptic plasticity.

For example, Sir John Eccles, recipient of the Nobel Prize for his work on synaptic transmission, and one of the legends of modern neuroscience, found that repetitive stimulation of nerves going to the spinal cord led to an increase in the size of the electrical response elicited in postsynaptic neurons in the spinal cord. Electrical stimulation of pathways was used as a simplified approach for activating synapses in a specific area of the nervous system. Though artificial, this method was viewed as a reasonable approximation of direct experience because our experiences in the world can affect the brain only by way of electrical conduction in nerve pathways (chap. 3). As a result, changes like those found by Eccles were interpreted as a step toward understanding the neural basis of learning. However, because the changes

mostly short-lived, they were not judged to be sufficient to account for the persistence of memory. Nevertheless, Eccles remained a true believer in a theory that memory involves synapses.

In 1966, Richard Thompson and Alden Spencer found indirect evidence that synaptic changes might account for a relatively simple example of learning. They studied habituation of the limb-withdrawal reflex in cats. Habituation is a form of learning in which repeated presentation of a stimulus leads to a weakening of a response—you jump the first time you hear a loud noise, but if it is repeated over and over, you jump less. The limb-withdrawal reflex had been known since Sir Charles Sherrington's time to be mediated by synaptic pathways into and out of the spinal cord (chap. 3). Later studies indicated that the actual circuit involved the relay of sensory information from the skin to a group of interneurons in the spinal cord to the motor (or output) neurons that control muscle movement. Thompson and Spencer ruled out changes in the ability of the input and output nerves to transmit signals, and concluded that the plastic changes must have crucially involved the interneurons. This important series of studies helped boost the notion that learning involves changes in synaptic transmission, but did not actually pinpoint the exact synaptic connections that were modified.

Two years later, Spencer and Eric Kandel took an important conceptual step, one that helped close the gap left open by Thompson and Spencer—the gap between behavioral learning and synapses. In their cellular-connection paper (chap. 5), they proposed that changes in synapses induced by learning could be identified if a simple behavior was studied in an animal with a simple nervous system. Cats, rats, or other mammals, whose behavior was more

relevant to humans, simply had too many neurons and too many synapses to be studied effectively. Lower vertebrates or, better yet, invertebrates were in their view more suitable subjects.

Kandel and other researchers followed this approach and went on to identify the synaptic basis of several forms of learning in the nervous system of invertebrate species. Especially notable was their ability to pursue synaptic plasticity all the way down to the level of specific molecules required to make memory last. Though truly groundbreaking from the point of view of a biological analysis of a behaviorally relevant form of neural plasticity, it was long unclear how, if at all, this work on lowly creatures like snails might apply to mammals.

In the meantime, though, a means of studying synaptic plasticity in the mammalian brain emerged. As this approach matured over the years, its findings began to converge with those from the invertebrate work, leading to the conclusion that, deep down, synaptic plasticity may be accomplished in similar ways in vastly different animals and in vastly different kinds of learning situations. After we take a look at the story of how it finally became possible to investigate synaptic plasticity in the mammalian brain, we'll return to the invertebrate studies to put the research in perspective.

PRACTICAL MAGIC

In the mid-1960s, Terje Lomo, working on his Ph.D. in Per Andersen's lab in Oslo, made a chance observation. He noticed that a brief burst of electrical stimuli delivered to fibers headed for the hippocampus in a rabbit led to a dramatic and long-lasting increase in transmission (a bigger electrical response to a test stimulus after, as compared to before, the burst) at synapses in the hippocampus, a region believed to be involved in human memory (chap. 5). Although this turns out to have been one of the most significant experimental findings in the history of memory research, when Lomo reported the result in a lecture at a scientific conference in 1966, there was, as he recalls, no reaction.

Lomo never got around to publishing his results and moved on to subjects of more interest to him. But a few years later, Tim Bliss, a young researcher from England, came to Oslo to work with Andersen. Bliss had been a student at McGill University, where he had attended Hebb's seminars, and he was quite interested in the physiology of memory. In Andersen's lab, Bliss and Lomo pursued the effects of electrical stimulation on hippocampal synaptic transmission that Lomo had noticed earlier, and in 1973, they published a paper describing the phenomenon of long-lasting potentiation, which is now called long-term potentiation (LTP).

In Bliss and Lomo's experiment, a stimulating electrode was put in the fiber pathway going into the hippocampus, and a recording electrode in the hippocampus itself. They then delivered a single electrical stimulus to the pathway, and recorded the electrical response of the postsynaptic neurons. This served as the baseline, the standard against which the rest of the experiment was gauged. Next they gave the potentiating stimulus—a brief burst of many rapidly repeated pulses. Then they started testing again with a single pulse, and continued testing periodically for several hours.

The key finding was that, following the potentiating pulses, the synaptic response got bigger, relative to the baseline response, and remained bigger for hours. According to Bliss, "We knew by the end of the first night that we really had something pretty important.

Although Bliss and Lomo had produced the long-lasting change in the postsynaptic response by electrically stimulating neural pathways, rather than by having the animals actually learn something, they realized that they had identified a mechanism that might be able to translate neural activity generated by environmental stimuli into changes in synaptic efficiency, a mechanism that might be used to record and store information about life's experiences. And the fact that they made this discovery in the hippocampus bolstered their concluding speculation that changes in the efficiency of synaptic transmission might account for memory.

It's interesting to note that the approach used by Bliss and Lomo to produce long-lasting changes in synaptic

transmission was very similar to the approach that had failed in the hands of Eccles and others in the spinal cord. Why was it successful with the hippocampus? One possibility is that the hippocampus is especially plastic. While this is true to a certain extent, LTP has since been induced in many areas of the nervous system, including the spinal cord. More likely, LTP was observed in the hippocampus because this part of the brain is organized into neat layers—the inputs come into one layer, and the outputs exit from a different one. This natural segregation made it possible to easily isolate and measure the postsynaptic responses elicited by electrical stimulation of input fibers. Once the phenomenon of LTP was worked out in the hippocampus, researchers knew what to look for and how to look for it in other neural circuits.

In the end, it took a lot of luck to achieve a viable model of plasticity in the mammalian brain. If the hippocampus didn't have such a simple organization, Lomo might not have been able to accidentally discover LTP. And if Bliss hadn't attended Hebb's seminars, or if he hadn't gone to Norway, Lomo's early finding might have remained in a drawer. But somehow the pieces all fell in place, and LTP became a way to practice Hebb's magic.

A SLICE OF LIFE

Bliss and Lomo did their studies in living animals. This was technically challenging, which is why the studies lasted throughout the night. But soon it was discovered that LTP could be studied more easily in thin slices of the hippocampus removed from the brain and submerged in a saltwater bath, with electrodes placed in the relevant regions. Hordes of scientists flocked to studies of hippocampal LTP. In the five-year period beginning in 1975, there were only twelve publications on LTP, whereas in the next five years, after the hippocampal slice technique became available, there were about ninety. But this was just the beginning. Between 1990 and 1994, more than one thousand LTP papers were published, and the number almost doubled in the next five years. The hippocampal slice became one of the preferred ways of studying regular synaptic transmission, as well as the premier way to study synaptic plasticity.

As the properties of LTP unfolded over the years, the possibility that this artificial phenomenon might have something to do with memory only increased. For example, in addition to its rapid induction and persistence over time, LTP was found to involve associative interactions between postsynaptic neurons and the specific presynaptic inputs that were involved in forming the association. Rapid acquisition, persistence, specificity, and associativity are all features one would expect of a memory mechanism. Let's consider the meaning of specificity and associativity further.

Bliss and Lomo showed that LTP was *specific* to the stimulated pathway. They started an experiment by stimulating two different pathways that elicit activity from the same population of postsynaptic neurons. They then gave only one of the pathways the potentiating stimulation, and subsequently tested postsynaptic response in both the pathways. Although stimulation of either pathway elicited a response from the postsynaptic cell, only the potentiated pathway changed. Giving potentiating inputs to one pathway, therefore, does not automatically change all of the synapses on the postsynaptic neuron; it only changes those synapses that were stimulated. LTP is thus specific to the synapses involved in the potentiating experience and does not represent a change throughout the entire postsynaptic neuron. This means that a given cell can participate in the storage of information about many different experiences so long as different synapses on the cell are involved in receiving them.

An important prelude to the demonstration of associativity was the finding that LTP involved cooperativity, which was shown by Bruce McNaughton and colleagues. Like Bliss and Lomo, they used two pathways. When each was stimulated in succession with weak stimuli, no LTP resulted. But if the

- two pathways were stimulated simultaneously with the weak stimuli, the stimuli could combine (cooperate) to produce LTP in both pathways. This suggested that LTP involves some kind of interaction at the cellular level between synaptic inputs, and might therefore be capable of forming associations between different inputs.

The ability to form associations between stimuli is perhaps the benchmark test for a synaptic mechanism of learning. That LTP might be a way to form associations was strongly suggested by the results of a study performed by Chip Levy and Oswald Steward in 1979. They applied weak stimulation to one pathway at the same time that strong

stimulation was delivered to another. In contrast to the McNaughton cooperativity experiment, in the Levy and Steward associative experiment, the strong stimulus alone was sufficient to induce LTP, and the weak stimulus added little. However, if the weak input arrived while the strong input was activating the postsynaptic cells, LTP occurred in the weak pathway as well as the strong one. Just as Hebb had predicted, if weak inputs arrive while the postsynaptic cell is active, the connection between the weak input pathway and the postsynaptic cell will be strengthened.

In 1986, several different research groups reported that the response of a single postsynaptic cell to a weak input could be potentiated if the post-synaptic cell was tricked into acting as though it had received a strong input. Using sophisticated techniques, the researchers were able to artificially reduce the electrical negativity of the cell just before a weak input arrived. (Recall from chapter 3 that when a cell is strongly activated by synaptic inputs, its internal electrical state becomes less negative, which is how an action potential is produced.) When a weak presynaptic input arrived during this induced

tivity, the synapse over which it was transmitted to the postsynaptic cell was strengthened. The combination of presynaptic and postsynaptic activity is indeed the magic formula for increasing synaptic strength.

ISN'T THAT A COINCIDENCE?

How does the brain actually achieve Hebbian plasticity? How, in other words, is the co-activity of presynaptic and postsynaptic cells registered and stored? Two discoveries in the mid-1980s began to clarify this process. The first was by Graham Collingridge, who showed that blockade of a particular type of glutamate receptor (chap. 3) prevented the induction of LTP without interfering with synaptic transmission. Thus, when this receptor was blocked, the synapse still worked just fine (release of transmitter from the presynaptic neuron produced a normal postsynaptic response), but it could not be potentiated by experience. The second discovery was by Gary Lynch and Roger Nicoll, who separately demonstrated that LTP would not occur if calcium was prevented from rising in the postsynaptic cell during action potentials. The two findings actually complement each other, since it is the special glutamate receptor that allows calcium to rise in the postsynaptic cell when an action potential occurs.

You'll recall from chapter 3 that glutamate is the main excitatory transmitter in the brain. When it is released from presynaptic terminals and binds to postsynaptic receptors, the likelihood of the postsynaptic cell's firing is increased. Actually, there are several different kinds of glutamate receptors, and each plays a different role. One (the AMPA receptor) is involved in regular synaptic transmission, while another (the NMDA receptor) is involved in synaptic plasticity. There are other glutamate receptors, but these two are the ones most relevant to this discussion.

Presynaptically released glutamate finds its way to both AMPA and NMDA receptors. Binding of glutamate to AMPA receptors is one of the major ways that a postsynaptic cell can be induced to fire an action potential, and is the means by which cells normally get fired up. In contrast, when presynaptically released glutamate reaches NMDA receptor on the postsynaptic cell, it has no effect initially because part of the receptor is blocked. However, once glutamate has activated the postsynaptic cell (caused it to fire an action potential) by binding to AMPA receptors, the block on the NMDA receptors is removed, and glutamate can open the receptor channel and allow calcium to enter the cell. LTP is the result.

For NMDA receptors to pass calcium, both presynaptic and postsynaptic cells must be active. This is the basic requirement for Hebbian plasticity. But how does this set of events form an association between two inputs to a cell? In the next section, we'll explore the chemistry of *how* calcium entry through NMDA receptors produces LTP. For now, though, what we want to consider is *why* calcium entry through NMDA receptors is a means of forming associations between a weak and a strong input.

Activity in the weak input pathway results in the release of glutamate and the binding of glutamate to postsynaptic receptors. Because the connection is weak, though, the input is not capable on its own of making the postsynaptic cell fire an action potential. However, when synaptic activity in the strong pathway activates the postsynaptic cell, the block on NMDA receptors will be removed, even at the weak synapses. Therefore, if the weak pathway is releasing glutamate during this time, NMDA receptors at both the strong and the weak synapses will be

able to bind the glutamate. Calcium will flow in through the NMDA receptors, and the weak synapses will be strengthened.

In sum, the reason that NMDA receptors allow LTP to occur is that they are coincidence detectors: they are able to register that presynaptic and post-synaptic neurons were active at the same time. More specifically, NMDA receptors allow the cell to record exactly which presynaptic inputs were active when the postsynaptic cell was firing. This input specificity is key to associativity, and is exactly what Hebb described decades before NMDA receptors were discovered. It is with this in mind that Holger Husi and Seth Grant recently referred to the NMDA receptor and associated molecules as a Hebbosome, a complex set of interacting proteins functioning as a unit in the induction and maintenance of synaptic plasticity.

MAKING CHANGE LAST

The binding of glutamate to its receptors is a brief event, lasting, at most, seconds. But memories can last a lifetime. In order for synaptic changes encoded by NMDA activity to persist, chemical processes that outlast the synaptic action itself are required. These processes have been studied using two different procedures, one that produces a form of Hebbian LTP that lasts only an hour or so, and another that generates a more persistent form of Hebbian LTP. These are called early and late LTP.³⁵ (There are also forms of LTP in which NMDA receptors play no role, but we will not have much to say about these.)

Early and late LTP are thought of as analogs of short- and long-term

memory. From a biological point of view, short- and long-term memory are also distinguished by their chemical requirements, in addition to their longevity. It has been known for several decades that if animals are given drugs that prevent the brain from making new proteins, they are able to learn normally but are unable to form long-lasting memories. (If they are tested within an hour or so of learning some task, they perform well, but show no signs of having learned the task when tested the next day.) This has turned out to be true of most if not all kinds of memory tasks that have been studied, and seems to apply to

most if not all species, as well. It is also true of early and late LTP: blockade of protein synthesis has no effect on early LTP but prevents late LTP. These parallels between early LTP and short-term memory; on the one hand, and late LTP and long-term memory, on the other, are consistent with the view that LTP and memory might be mediated by the same molecular mechanisms.

In order to further understand the chemistry of early and late LTP, and to pursue the question of whether LTP and memory have similar molecular underpinnings, we need to examine the concept of second messengers. Neurotransmitters like glutamate are considered first messengers. They are responsible for signaling between neurons. Second messengers pick up where first messengers leave off. Their job is to initiate chemical reactions *within* the cell on the basis of information provided from *outside* the cell during neurotransmission by first messengers.

Calcium is one of the major second messengers. As we've seen, when glutamate binds to its NMDA receptors, calcium flows into a cell (after the block has been removed). Once this occurs, calcium takes over and directs the chemical reactions that strengthen synaptic connections, both in the short run and the long run.

Key to the whole process are enzymes called protein kinases that activate specific proteins. Their job is to phosphorylate certain proteins. Technically, this means that they add a phosphate group to the proteins. But all you really need to know is that phosphorylation turns proteins on, transforming them from an inactive to an active state.

In early LTP, kinases act on preexisting proteins that are present in the cell, waiting to be called upon. As a result of calcium influx during early LTP, several kinases are activated.. One important task performed by these kinases appears to be the phosphorylation of AMPA receptor proteins. As a result, after the induction of LTP, the same amount of glutamate released from the presynaptic cell by an action potential will have more AMPA receptors to bind to, and a bigger postsynaptic response will occur, allowing each presynaptic action potential to make a greater contribution toward the firing of the postsynaptic cell. (Recall from chapter 3 that it takes many closely timed action

potentials at different synapses on a postsynaptic cell to generate an action potential in that cell, so every little bit helps.)

Although most scientists in this area of research agree that LTP is induced or initiated by calcium-triggered chemical reactions in the postsynaptic cell, some believe that these postsynaptic changes fully account for LTP, while others argue that the presynaptic cell changes as well. For example, a popular notion is that after LTP has been induced, the presynaptic cell more easily releases glutamate when an action potential occurs in its terminal, leading to a bigger postsynaptic response. Evidence for this theory has been obtained in sophisticated experiments, involving the precise measurement of glutamate release from presynaptic neurons, performed by Charles Stevens of the Salk Institute and Richard Tsien of Stanford, among others.” However, because LTP is triggered only in the postsynaptic cell, in order for the presynaptic hypothesis to be valid, something has to be communicated from the postsynaptic to the presynaptic neuron after LTP is initiated. One possible way this could happen is by means of retrograde messengers, substances released by the postsynaptic cell after LTP is induced and taken up by the presynaptic terminal, which cause changes in the ease of glutamate release there. Robert Hawkins in Kandel’s group and others have found evidence that retrograde messengers do modify the presynaptic terminal in ways suggested by the presynaptic hypothesis, but as other leading LTP researchers, such as Rob Malenka and Roger Nicoll, point out, this notion remains controversial.

While early LTP involves the activation of existing proteins by kinases, the creation of long-lasting or late LTP, like long-term memory, involves the formation of new proteins. One of the key steps in this process is the activation of kinases, especially protein kinase A (PKA), MAP kinase (MAPK), and calcium/calmodulin protein kinase (CaMK), in a special way that allows them to move inside the cell nucleus. Once there, they activate a protein called CREB. The job of this gene transcription factor is to activate specific genes that make new proteins that travel back to the synapses that started the whole process and stabilize the connections (see “Tagging Along,” below).

Our understanding of the molecular basis of LTP has come from two kinds of studies. The traditional approach has involved the use of drugs that block certain molecular steps, like the flow of calcium through NMDA receptors, or phosphorylation of kinases, or synthesis of proteins. In studies of this type, the drugs are either placed in the bath in which the slice of brain is maintained or are injected directly into the postsynaptic cell. If LTP is disrupted, the chemical step interfered with is implicated in the underlying plasticity. Recently, though, a new approach has emerged. Rather than adding drugs to a brain slice, mice are created that either lack or have an excess of specific molecules, like CREB or certain kinases, as a result of genetic alterations. LTP is then induced in brain slices taken from these mice. As in the drug studies, if LTP is affected, then the missing or exaggerated molecule is implicated.

Most work to date using genetically altered mice has involved permanent deletion of a gene in a line of animals, leading to the absence of the particular molecule in the entire body throughout life. The fact that the molecule is missing throughout life, however, raises some questions about how to interpret the data, since an effect on LTP could be due to the absence of the molecule itself, or could be a secondary effect of the animal’s having gone through life without the molecule. The value of this approach is being refined by the ability to selectively alter molecules in specific brain regions, like the hippocampus, leaving other areas intact, and by the ability to turn certain genes on or off at a certain time. As a result, the mouse can grow up with all its molecules intact, and then right before the LTP experiment a chemical switch can be thrown to either decrease or increase the molecule being studied. While the latter approach is obviously more powerful, it is also more difficult to implement. But researchers are working diligently to improve these techniques, hoping to be able to turn genes on or off in specific brain areas on command. Genetically altered mice have already played key roles in brain research, implicating NMDA receptors, protein kinases, and CREB in LTP, and are likely to become even more valuable as the techniques for creating genetically altered organisms improve in the coming years.

I’ve mentioned only briefly some of the many molecules implicated in Hebbian LTP. Our knowledge of the role of these substances is expanding rapidly. And while there remain many unresolved mysteries about the molecular basis of synaptic plasticity, at a minimum it seems fairly clear that NMDA receptors, calcium, certain kinases, CREB-activated genes, and protein synthesis are involved. The massive amount of data pouring forth on the molecular basis of LTP is a bit overwhelming, so much so that Josh Sanes and Jeff Lichtman, two neuroscientists who work on brain

development, have called for a moratorium on molecular studies of LTP. Others, like David Sweatt and Mary Kennedy, see it as a sign of the vitality of the field that so much information is being generated about how LTP works. I share their optimism.

TAGGING ALONG

The synaptic activity that triggers LTP occurs out on the dendrites, while the genes that make the proteins that enable LTP to persist are located far away in the cell nucleus. We've already seen how the gap from the synapse to the nucleus is spanned: calcium comes in through the NMDA receptor and triggers molecular changes that ultimately affect genes in the nucleus. More mysterious has been the way this circle is closed. How, in other words, do proteins made in the cell body know which of the many synapses on the neuron originally were responsible for the initiation of the protein synthesis? It's these synapses, and these alone, that have to be modified in order for the plasticity to be specific to the active synapses.

Recent studies have helped solve some of the mystery. It appears from work by two research groups (Eric Kandel and Kelsey Martin in New York, and Richard Morris and Uwe Frey in Scotland and Germany) that, following some significant experience, the active synapses are given a molecular tag. When the new proteins are subsequently made in the cell body, they are then shipped out to essentially all possible synapses in the cell, but only synapses that were tagged during the initial stimulation are able to use the new proteins to stabilize the connection with the presynaptic terminal.

Once they find their way back to the tagged synapses, the proteins stabilize the connection between the presynaptic and postsynaptic neuron. There are two parts to the explanation of how this occurs. The first has to do with the stabilization of existing synapses. These are kept in the new, improved, facilitated state after LTP in several different ways. For example, by making more AMPA receptors available on a prolonged basis, a given amount of glutamate released by the presynaptic terminal continues to have a bigger effect. Fur-

ther, as we saw, terminals connected to the postsynaptic cell may come to release transmitter more efficiently as a result of a retrograde messenger passed back from the postsynaptic cell to the presynaptic terminal. By prolonging temporary changes, the effects of LTP are made to endure.

The other part of the explanation is related to the actual growth of new synapses. Several researchers have shown that LTP does indeed create new connections. These appear to be fostered in part by the release of growth-inducing tonics, neurotrophins, from the postsynaptic cell. We discussed a similar role for neurotrophins in the context of development in chapter 4. Neurotrophins are released from the postsynaptic cell when it is active. They are then taken up by the presynaptic terminals that were also just active, and encourage the terminal to sprout new branches and form new synaptic connections with the postsynaptic neurons. The availability of more synaptic connections means that a given action potential coming down an axon will have a bigger effect, since it has more places (terminal branches) to release transmitter from and more postsynaptic sites to bind to.

Recent studies demonstrate that proteins can be synthesized in dendrites? If dendritic protein synthesis occurs during memory formation it would simplify, to some extent, the problem of achieving synapse specificity since proteins made locally would not have to find their way back from the cell body to the active dendrites. This is an exciting new area of research.

FROM LTP TO MEMORY

In the end, what do LTP, NMDA receptors, kinases, and CREB have to do with memory, as opposed to brain cells living in a dish and being stimulated electrically? Let's start with NMDA receptors.

The first bit of evidence supporting a role for NMDA receptors in memory came from a study performed by Richard Morris, working in collaboration with Gary Lynch in Irvine, California. Morris used a drug that specifically blocks NMDA receptors, leaving AN4PA receptors and other aspects of brain function unaffected. He put this drug into

the brain of rats in a way that bathed the hippocampus, and then tested the rats in his famous water maze—one of the better tasks for implicating the rat hippocampus in spatial memory (chap. 5). Control rats were able to learn where the submerged platform was located. Rats treated with the drug, though, had great difficulty learning this task in spite of the fact that they were able to swim to the platform if it was above water, showing that the drug did not affect normal vision or swimming ability. Subsequently, when the brain was removed from these rats and slices of the hippocampus kept alive in a dish, LTP could be readily induced in the controls, but LTP induction was impaired in the drug-treated brains. Since hippocampal LTP and hippocampal memory were both disrupted by the same treatment, the conclusion was that hippocampal LTP might have something to do with hippocampal memory.

Many subsequent studies have followed up on these important observations. There have been successes and failures, as well as alternate interpretations, in the quest to relate NMDA receptors in the hippocampus to spatial memory. Particularly difficult to ascertain has been the issue of whether NMDA blockade produces a specific effect on learning itself, or instead changes other less specific capacities, such as the ability to perceive a given stimulus or the ability or desire to perform the appropriate response. However, the most recent work by Morris and colleagues seems to show fairly convincingly that blockade of NMDA receptors in the hippocampus does interfere with spatial learning.

That hippocampal NMDA receptors are involved in spatial learning is also suggested by the results of studies using genetically altered mice. Susumu Tonegawa and colleagues at MIT created mice lacking important components of NMDA receptors. Particularly significant was the fact that they were able to restrict the loss to the hippocampus—in fact, to a specific part of the hippocampus. The mice performed poorly on spatial learning tasks, and LTP induction was impaired as well. More recently, Joe Tsien, a former colleague of Tonegawa's, took the opposite approach, enhancing the function of NMDA receptors in a specific region of the hippocampus. LTP induction was facilitated, and so, too, were several memory functions dependent on the hippocampus. These studies provide fairly direct evidence for the three-way relation between hippocampal NMDA receptors, hippocampal LTP, and hippocampal-dependent memory.

Similar studies comparing LTP and spatial memory have also been performed for some of the other molecules that have been implicated in LTP. The literature is at this point fairly extensive but incomplete, so I'll mention only a few findings. Alcino Silva and colleagues examined the effects of knocking out CREB in mice, and discovered disruption in both LTP and hippocampal-dependent memory. Further, Ted Abel, Mark Mayford, Eric Kandel, and colleagues have found that genetic alteration of protein kinase A and a calcium/calmodulin-dependent kinase, a major enzyme that phosphorylates CREB, disrupts late LTP and long-term memory without affecting early LTP or short-term memory.'

Considerable research thus supports the view that hippocampal LTP and hippocampal-dependent memory operate by similar molecular mechanisms. The parallels, in fact, are quite impressive, but not perfect. For example, in some studies, hippocampal-dependent learning is unaffected by treatments that block LTP. This type of result, though, has been fairly rare, and certainly is not compelling enough to undermine the whole field. Further, we now know that there are a variety of different forms of LTP, some involving NMDA receptors and some not, and some Hebbian (requiring presynaptic and postsynaptic activity) and some not (possibly requiring changes only presynaptically), and there's even long-term depression, where synapses are weakened by uncorrelated activity between presynaptic and postsynaptic cells. While we have a long way to go before we fully understand how the synapses make memories, at a minimum, LTP has been an excellent tool for pursuing the topic.

A MISSING LINK

Underlying much of the work on LTP is the assumption that LTP is not just a way to study how experience changes synapses, but is, in fact, *the* way that synapses are changed when we learn. The various findings indicating that the same molecules are involved in LTP and memory are consistent with the view that LTP occurs during learning, but this evidence is, as critics have noted, circumstantial! In an effort to silence these critics, LTP researchers have attempted to demonstrate that something like LTP does take place in the hippocampus during

learning. However, these studies have, for various reasons, come up short. For a long time, I myself was ambivalent about LTP. I thought it was an interesting phenomenon but was not convinced that it was the answer to how memories are created. I became more intrigued when Marie-Christine Clugnet and I found that we could induce LTP in a neural pathway carrying auditory information to the amygdala, a pathway my lab had been studying because of its role in fear conditioning. But what finally made me a convert were studies performed in my lab by Michael Rogan while he was working toward his PhD. Now I am a proselytizer.

Rogan decided to take on the challenge of trying to show definitively that LTP occurs during learning. While most of the earlier work in this area had involved the hippocampus, the lack of understanding of the actual circuits involved in hippocampal-dependent memory (spatial memory) was a barrier. Rogan instead started with a circuit known to be involved in fear conditioning and then asked whether LTP occurred in that circuit. This reversal of the strategy (starting with a learning circuit and looking for plasticity in it, rather than starting with a form of plasticity and asking how it might relate to learning) turned out to be very advantageous.

My lab had been studying the anatomy of fear conditioning in rats and had shown that in order for a sound to acquire aversive properties when associated with a shock, the sound had to be relayed from the auditory thalamus to the lateral nucleus of the amygdala. And, as I noted above, we had also found that we could induce LTP in this same pathway. Rogan put these findings together, and then took two additional steps.

First, given that we used sound stimuli in conditioning studies, and that sounds get to the amygdala through the thalamoamygdala pathway, Rogan asked whether induction of LTP in this pathway would change the way sound stimuli are processed in the amygdala. He induced LTP in the usual manner with rapid electrical stimuli, but made a key change in the way LTP was tested—using a natural stimulus (in this case, a sound) rather than electrical stimulation of nerve fibers. His findings showed that induction of LTP in the pathway that carries the sound to the amygdala enhanced the amygdala's response to that sound. When this article was published, it was accompanied by commentaries from several leaders in the field who noted that this was indeed an important step forward in the long quest to relate LTP and memory:

Rogan's first study showed that artificial alteration of synaptic transmission in the pathway changes the way that pathway processes external stimuli. This was an interesting discovery, but it left unanswered the question of whether, during natural learning, something like LTP occurs in the brain. In the second study, he therefore substituted fear conditioning for LTP induction. A few years earlier, Liz Romanski had shown that single cells in the lateral amygdala receive both sound and shock information, thus demonstrating that they might be coincidence detectors. It thus seemed reasonable that fear conditioning with a shock might change the responses elicited from lateral amygdala cells by sound stimuli. Indeed, Rogan's results showed that fear conditioning and LTP induction produced very similar changes in the electrical responses of amygdala cells to sound stimuli. Fear conditioning, in other words, seems to induce LTP. This work, which earned Rogan the prestigious Lindsley Prize in 1999 for the best dissertation in behavioral neuroscience, prompted a commentary by Charles Stevens, one of the top researchers in LTP, who asked whether the "million-dollar question" about

memory had finally been answered. He concluded that though it had not, we had taken a key step toward doing so.

Fear conditioning has thus been more successful in closing the gap between LTP and memory than studies of spatial learning or other hippocampal-dependent tasks. On the basis of careful anatomical and behavioral studies, we had developed a clear idea about which synapses might be changing during fear conditioning, and we could easily stimulate them with both natural and electrical stimuli. While the circuitry of the hippocampus is well understood, the relation of learning to specific circuits in the hippocampus has been difficult to determine because of the complexity of the learning tasks. In spatial learning, a rat is free to learn about any of the many stimuli present in the environment, making it difficult to determine, in Hebbian terms, precisely what's being associated with what. In contrast, fear conditioning is a straightforward form of associative learning that is easy to define in terms of Hebbian plasticity, where strong (foot shock) and weak (sound) stimuli interact at synapses on the same cells.

An important question is whether the huge amount of work that has been done on the molecular basis of hippocampal LTP applies equally to amygdala LTP and fear conditioning. Fortunately, much of this research does seem relevant. For example, studies by several researchers, including Sarina Rodrigues in my lab, and Mike Davis and colleagues and Mike Fanselow and colleagues (especially Jeansok Kim and Steve Maren) have shown that blockade of NMDA receptors in the lateral amygdala prevents the acquisition of fear conditioning. Further, Glenn Schafe in my lab has shown that disruption in the lateral amygdala of protein synthesis or of some of the same kinases that prevent the induction of late LTP in the hippocampus interferes with the formation of long-term memory for fear conditioning without affecting short-term memory. Some of the main kinases involved are PKA and MAP kinase. PKA has also been implicated in context conditioning by Schafe and by Rusiko Bourtchouladze in Kandel's lab, but the exact site of action in the brain is not known. Also, when CREB levels are increased in the region of the lateral amygdala, weak training produces strong learning. CREB, PKA, and other molecules have also been implicated in aspects of fear conditioning through studies of genetically altered mice by Eric Kandel, Alcion Silva, Mark Mayford, Tsunmo Tonegawa, Joe Tsien, and others. Importantly, Y. Y Huang and Kandel have shown that late-phase LTP in the lateral amygdala also depends in part on NMDA receptors, PKA, MAP kinase, and protein synthesis, further closing the gap between memory formation and LTP in the lateral amygdala. Fear conditioning and lateral amygdala LTP thus seem to require the same mix of molecules that has been implicated in synaptic plasticity and memory in the hippocampus: Hebbian changes in synaptic strength, which are triggered by calcium entry through NMDA receptors, which are stabilized by protein kinases that induce CREB-related genes to make new proteins. NMDA receptors, MAP kinase, CREB, and protein synthesis also have been implicated in classical conditioning, by Yadin Dudai and colleagues at the Weizmann Institute in Israel (especially Raffi Lamprecht, Kobi Rosenblum, and Diego Berman).

What is the significance of the fact that fear conditioning and taste-aversion learning by the amygdala and relational (spatial) learning by the hippocampus all involve similar molecular changes? Obviously, the circuits are very different at several levels. For example, the fear-conditioning circuits of the amygdala are one synapse removed from sensory systems, while the hippocampus is several synaptic steps removed from them. In addition, the consequences of activating these circuits are quite distinct. Activation of fear-conditioning circuits in the amygdala leads to the expression of hard-wired bodily responses by way of connections with the brain stem, whereas activation of hippocampal circuits leads to representations that allow for a multitude of responses through widespread connections with cortical areas. Different forms of memory are thus distinguished not by the molecules that make them but by the circuits in which those molecules act.

There is still quite a lot of work to be done on the molecular basis of fear conditioning to pin down some of the particulars. For example, recent studies in brain slices of the amygdala by Marc Weisskopf and Liz Bauer in my lab have suggested that at least some of the plasticity may be mediated by calcium entry through special calcium channels. Following up on this work, Bauer and Glenn Schafe showed that fear conditioning could be prevented by blocking these calcium channels, suggesting that calcium entry through both NMDA receptors and calcium channels may contribute to fear learning. The theory is being evaluated in ongoing studies. In hippocampal LTP, this combined involvement of NMDA receptors and calcium channels has been seen. Other research indicates that presynaptic plasticity may occur at some synapses in the amygdala during LTP induction and during fear conditioning². If so, classical fear conditioning in the rat might turn out to be similar to classical conditioning in invertebrates, in that both presynaptic and postsynaptic changes take place. We'll revisit this issue below when we consider invertebrate learning. Also unresolved is the locus of memory storage— plasticity in the lateral amygdala is clearly involved, but whether this is the only site of change that is maintained over time remains to be determined.

Studies of fear conditioning have also revived interest in a strange but possibly very significant phenomenon in memory research—reconsolidation. The recent discovery, made by Karim Nader and Glenn Schafe in my lab, is that protein synthesis in the amygdala seems necessary for a recently activated memory to be kept as a memory. That is, if you take a memory out of storage you have to make new proteins (you have to restore, or reconsolidate it)

in order for the memory to remain a memory. One way of thinking about this is that the brain that does the remembering is not the brain that formed the initial memory. In order for the old memory to make sense in the current brain, the memory has to be updated. This work has stimulated a lot of interest from both scientists and lay persons. One man called and asked whether it might be possible for him to eliminate the memory of his ex-wife by blocking protein synthesis in his brain while thinking of her. The practical side of this is that it might be possible some day to have trauma victims recall conditioned taste-aversion learning, another form of aversive their trauma in the presence of some drug or other brain alteration that reduces the stranglehold of the memory on the person's psyche. After we proposed this, though, a therapist made a very good point. What would it mean to a Holocaust survivor, for example, to lose such memories after having lived for many years and having developed an identity based in part on them? This is a very important concern, and touches on the deep ethical issues that scientific discoveries can raise.

Fear-conditioning research has capitalized on breakthroughs in modern neuroscience to move us forward in our understanding of how synaptic changes underlie memory. Consistent with the cellular-connection strategy outlined by Kandel and Spencer in 1968, a simple behavior was chosen and the circuit was mapped. The plasticity of cells in the circuit was studied. The cellular changes have been related to specific synapses, and the synaptic changes have been related to specific molecular events. It was inconceivable in 1968 that this could be done in a mammalian brain, and it is a testimony to how fast the field of neuroscience has advanced that so much has been achieved in recent years.

SNAIL TALES

The molecules of memory are conserved not just across different kinds of memory within a species but also across vastly different kinds of organisms. In fact, many of the basic facts about the molecular basis of plasticity were first revealed in studies of invertebrate organisms, and only later were they found to be applicable to plasticity in the mammalian brain through studies of LTP and behavioral learning. It is time to put the mammalian brain aside and take a look at some of the research on invertebrates that led to our current molecular understanding of memory.

The neural basis of learning and memory has been investigated in many different invertebrates (bees, grasshoppers, crayfish, slugs, flies, and various mollusks), but studies of the mollusk *Aplysia californica* have been particularly thorough and informative. Much of the work on this marine snail has been conducted by Eric Kandel and his students and colleagues at Columbia University and elsewhere. This pioneering research was a major factor in Kandel's receipt of the Nobel Prize in 2000.

BEHAVIORAL STUDIES

Aplysia breathe through gills that are covered by a piece of skin called the mantle. If the mantle is lightly touched, the gill retracts. This defensive reflex

protects the gill from injury, and has been the subject of extensive investigation as a behavioral model of learning and memory. The plasticity of other reflexes in *Aplysia* has also been studied, but the gill-withdrawal reflex will be used to illustrate the basic findings.

The central nervous system of an *Aplysia* has only about twenty thousand neurons. It has been estimated that fewer than one thousand neurons (and possibly as few as four hundred) are active during the gill-withdrawal reflex, and many of these are not believed to be essential. Essential ones include sensory neurons that process the touch of the mantle skin, and motor or output neurons that control the gill withdrawal. The sensory neurons make synaptic connections with the motor neurons. Also important are sets of excitatory and inhibitory interneurons that make synapses with the sensory or motor neurons and regulate the reflex. When you compare this with the billions of neurons in the mammalian brain, it is clear why it might be easier to investigate the neural basis of behavioral functions in the *Aplysia*.

The gill-withdrawal reflex exhibits several forms of learning, including habituation: the gill retracts less if the mantle is touched repeatedly. Habituation is a form of nonassociative learning: a single stimulus is involved, and it is not associated with anything else. Habituation can be reversed rapidly by giving a strong stimulus, such as an electric shock, to some other part of the mollusk's body, like the tail. Touch of the mantle after a shock results in a big response. This is an independent form of nonassociative learning called sensitization and is not simply a recovery from habituation, since the response to the touch of a second area of skin, one that was not habituated, is also amplified by the shock. Sensitization is considered nonassociative because the touch-test stimulus and the shock were never related, and the learning is not specific to the test stimulus—many stimuli given after the shock produce bigger effects. The effects of sensitization can be short-lived or long-lasting, depending on the strength of training. A weak shock given only once will produce short-lasting changes in the reflex that disappear within a matter of hours, whereas repeated presentation of the shock leads to changes that can last for days. Like habituation, sensitization is also regarded as a form of nonassociative learning because a single stimulus (the shock) causes the behavioral change.

Associative learning—namely, classical defense or fear conditioning—also occurs in the gill-withdrawal reflex. This is essentially the same kind of conditioning task that we've already encountered several times in mammals. See, for example, the discussion of implicit learning in chapter. If a shock is delivered to the tail while the mantle is touched, subsequently the touch alone will lead to a stronger withdrawal of the gill than it did before conditioning. That this is associative learning (and not just sensitization due to the shock) is demonstrated by the fact that the conditioned response is smaller if the touch and shock do not occur at the same time. The relation between the two stimuli is thus key. The associative nature of the learning is further indicated by the fact that if two touch stimuli are used (that is, if different parts of the mantle are touched), and only one is paired with the shock, the paired touch elicits a bigger response than the unpaired one.

Conditioning and sensitization are similar in that a strong stimulus changes the response to a weak one. But they differ in specificity: in associative conditioning, the amplified response only occurs in reaction to a stimulus that was paired with the shock, whereas in sensitization the response to stimuli that have no relation to the shock is also bigger. Sensitization basically makes the snail jumpy, so that any stimulus to which it is exposed after a strong stimulus will lead it to react, whereas it is conditioned if it reacts to a stimulus that occurred at the same time as a strong stimulus but not to other novel stimuli that occur afterward. Studies of classical conditioning often use unpaired training (sensitization) as a control to ensure that the conditioned response is due to the associative relation between the strong and weak stimulus, as opposed to the nonassociative effects of the strong stimulus alone.

The big advantage of studies of animals like the *Aplysia* is, of course, that the small number of neurons involved in the control of a behavioral response like the gill-withdrawal reflex makes it relatively easy to pursue questions about exactly which neurons and synapses are changed by experience, and what sorts of molecular events might underlie the synaptic alterations that constitute the memory. But even in a simple nervous system like that of the *Aplysia*, it has proven useful to narrow the focus down even further. Because the actual behavioral response is controlled by a number of different neurons, some studies sidestepped behavior completely and instead examined the electrical activity of a single motor neuron as an indicator of whether learning occurred. In this experimental setup, the electrical activity of motor neurons in response to natural stimuli (touches and shocks) or to electrical stimulation of the sensory pathways that transmit natural stimuli can be assessed. The sensory or motor neurons also can be directly treated with chemicals, like drugs that mimic or block the effects of neurotransmission or other aspects of cell function. A further reductive step has involved complete elimination of the *Aplysia*, except for a single sensory and a single motor neuron that grow synapses between each other in a dish. The response of the motor neuron to electrical stimulation of the sensory neuron or to direct application of drugs to the neuron can then be tested. With these approaches, *Aplysia* researchers have learned much about how synaptic changes such as those occurring during learning are established and maintained.

MEMORY MECHANISMS

All forms of gill-reflex learning involve changes in synapses between sensory neurons that receive inputs from the mantle skin and motor neurons that control the gill response. In habituation, for example, the response of the postsynaptic neuron to a presynaptic input weakens, and the gill response gets smaller, because the presynaptic terminal comes to release less glutamate. It simply gets depleted.

In contrast, in sensitization, the gill reacts more to the same stimulus after the tail is shocked than before because the sensory neuron comes to release more glutamate. To understand why more glutamate comes out after the tail is shocked requires that we consider the way the shock pathway interacts with the synapses that connect the sensory and motor neurons.

The shock pathway forms synapses on the terminals of the sensory neurons. These are called axoaxonic synapses, since the axon terminals of the shock pathway make synapses with other axon terminals, in this case, terminals of the sensory pathway. This contrasts with the situation we've considered most so far, in which axon terminals contact dendrites. The tail shock pathway thus ends on the presynaptic terminal of the sensory neuron and causes more transmitter to be released. Because the increase in the efficiency of transmission between the sensory and motor neuron is due entirely to alterations in the sensory neuron terminal, it is referred to as presynaptic facilitation. Because the state of activity of the postsynaptic neuron is irrelevant, sensitization by definition is a form of non-Hebbian plasticity (recall that Hebbian plasticity requires presynaptic and postsynaptic activity).

How does the tail shock cause presynaptic facilitation? The tail shock pathway releases the modulator serotonin at the sensory terminal. When serotonin binds to its receptors on the terminal, it activates second messengers there that ultimately lead to the activation of a protein kinase, namely PKA, which is also involved in LTP. PKA in turn leads to other changes that cause action potentials to last a little longer than usual. As a result, more glutamate is released from the terminal after the shock than before. This causes a bigger response in the motor neuron, and therefore behaviors controlled by motor neurons, like gill withdrawal, are more strongly expressed.

The sensitizing effects of shock are short-lived unless the shock is given repeatedly. When a repeated shock is administered, additional processes are activated, and the effects of sensitization can last for days. In particular, PKA is activated in a special way that allows it to enter the cell nucleus. In addition, a second protein kinase, MAP kinase (which is involved in hippocampal and amygdala plasticity), is also activated by repeated shocks and moves inside the nucleus. PKA and MAP kinase then phosphorylate the gene transcription factor, CREB, which we also encountered earlier in our survey of LTP.

CREB-activated genes make new proteins that end up facilitating transmission between the sensory and motor neuron in various ways. One effect of these proteins is that PKA becomes persistently activated, which means that the short-term sensitizing effects of PKA are lengthened, continuing the effects on transmitter release described above. Presynaptic facilitation simply continues. But there is at least one other effect that is also important: some of the proteins stimulate the sensory terminal to grow new axonal branches that make new synapses with the motor neuron. With more synaptic connections between two cells, a given action potential in the presynaptic axon has more ways to influence the postsynaptic cell, leading to a stronger response. The two effects (persistent activation of PKA and formation of new connections) combine to produce long-lasting increases in transmission between the sensory and motor neurons, and stronger reflex responses.

Classical conditioning, though obviously associative in nature, was long believed to be mediated by a non-Hebbian synaptic mechanism, namely, an amplification of presynaptic facilitation. That is, when synaptic activity occurred in the presynaptic terminal while the transmitter release was being facilitated (sensitized) by the shock, the degree of facilitation increased and came to be coupled to the stimulus that elicited the activity. This was considered associative, since it involved changes restricted to the stimulus that was activating the terminal, but it was non-Hebbian because the postsynaptic cell was not believed to be involved. However, recent work has shown that the initial studies on which this view was based were incomplete. It is now known that the postsynaptic cell does participate in associative classical conditioning. Researchers like David Glanzman and Jack Byrne argue that classical conditioning in the *Aplysia* does indeed involve Hebbian synaptic changes, and probably also involves

NMDA receptors, in addition to an activity-dependent enhancement of presynaptic facilitation. One way of understanding this is that nonassociative (non-Hebbian) presynaptic facilitation functions as the presynaptic component of associative (Hebbian) plasticity.

Considerable work has also been conducted on the molecular basis of classical conditioning in the *Aplysia*. While some of the early steps are different from those involved in sensitization, the long-term changes involve phosphorylation of CREB by protein kinases and synthesis of new proteins by CREB-activated genes. The establishment of long-term plasticity seems to follow a common pattern even when it is triggered by different kinds of short-term changes.

The most obvious differences between associative conditioning in the *Aplysia* and plasticity in the mammalian brain is thus the importance of presynaptic plasticity in the *Aplysia*. The recent finding that NMDA-mediated plasticity in the postsynaptic cell is also significant in *Aplysia* plasticity helps close the gap. And the possibility that fear conditioning, amygdala LTP, and several forms of hippocampal plasticity might, like plasticity in *Aplysia*, involve presynaptic as well as postsynaptic changes further strengthens the notion that similar mechanisms are used to make memories in diverse species, suggesting that memory mechanisms are conserved across many levels of evolution.

FLY GENES

Research on the fruit fly, *Drosophila melanogaster*, has also been instrumental in the accumulation of our knowledge of memory and its molecules. But in contrast to the highly focused cellular-connection approach used to study the molecular basis of plasticity in *Aplysia*, the fly work has focused on behavioral studies of mutant strains of flies that lack specific molecules, and has for the most part not related the molecular effects on learning and memory to specific cells and synapses.

The studies of genetically altered mice we discussed earlier are descendants of fly mutation studies started several decades earlier by Seymour Benzer and his colleagues. Although Benzer started investigating fly behavior in the

1960s, the use of flies in research on learning began in the following decade, when Chip Quinn and Yadin Dudai in Benzer's lab showed that normal flies could remember a smell associated with electric shock. When given the choice between going into a chamber that contained an odor associated with the shock and a chamber that held a different odor, the flies often avoided the former chamber.

Most work on mutant flies has involved this kind of learning task, which is simply a variation of classical defense or fear conditioning. These flies are called learning mutants, since they are altered genetically in a way that causes them to fail to learn or remember. Learning mutants are created by treating normal flies with a chemical that produces more or less random mutations in their DNA. The flies are then tested for deficits in learning and memory. Those that are deficient are bred following precise mating schemes that create lines of genetically identical flies that have learning and memory impairments.

Genetic mutations can disrupt learning for a variety of different reasons that have nothing to do directly with learning and memory, and it is necessary to rule these out when drawing conclusions about whether a mutant has a specific deficit in these areas. For example, a fly that cannot smell or that is defective in perceiving a shock, or one that is chronically ill, will not learn. These are concerns that apply to the studies of genetically altered mice described earlier, as well. In addition, as in the mouse studies, it is important to keep in mind that a mutant animal goes through its entire life with the genetic alteration, sometimes making it difficult to distinguish effects on learning and memory from other more general consequences of having grown up without certain proteins and their products. Behavioral changes in a mutant could, in other words, be due either to the mutation itself or to the consequences of surviving with the mutation.

The first fly learning mutant to be identified was called dunce. Then came amnesiac, cabbage, rutabaga, turnip. All of these variants performed poorly on the fear-conditioning task. In each of these mutants, learning actually

did take place to some extent, but then decayed rapidly after training. Memory that remains after the initial decay, however, is fairly stable. This suggests that the mutations affect the early more than the late phase of memory.

Most of the initial work on learning mutants focused on relatively shortlasting memories. However, recent studies by Tim Tully, Jerry Yin, and Chip Quinn showed that long-lasting memory that persists for days can be created by giving normal *Drosophila* multiple training trials that are spread out over-time (this is called spaced training). Disruption of protein synthesis prevents this long-term memory from forming but has no effect on a shorter-lasting kind of memory induced by a few training trials given all at once (called massed training). Like a college student, the fly does better if it takes in information at a regular pace over time than if it crams.

Yin and Tully made use of this difference in massed vs. spaced training to examine whether cAMP-mediated gene induction and protein synthesis underlies long-lasting memory. They employed genetic techniques (see above) to increase the level of the gene transcription factor CREB available in flies at the time just before training. This step overcomes many concerns about whether the deficit is due to a failure to form memory or to some irrelevant consequence of growing up without the gene. The researchers found that in the animals with extra CREB, a single training trial could accomplish what required multiple spaced trials in normal flies—create a memory that lasts for days. This suggests that training-induced CREB activation (and the consequent induction of CREB-related genes) is an essential step in the production of proteins essential for long-term memory. The molecules leading to induction of CREB-related genes and their proteins are thus the stars of memory across species and training procedures.

Flies are ideal in many ways for performing genetic studies of learning and memory. They exhibit a robust form of learning, and because of their brief life span, many generations can be studied in a short time. But they are not especially well-suited specimens for pursuing the individual cells and synapses involved in learning and memory. Although recent work has made some progress in pinpointing the locus of plasticity, circuit analysis in the fly will probably remain difficult due to the small size of the nervous system in this tiny creature.

I HAVE TO ADMIT IT'S GETTING BETTER

A key question you may now be asking yourself is whether all this hard-core neuroscience has, in fact, any practical application. In other words, might it be possible to use this kind of work to help improve normal memory and, even more important, to rescue or prevent age-related memory loss?

One of the pioneers in LTP research, Gary Lynch of the University of California, has been looking into ways of altering the molecular basis of synaptic plasticity in the hope of finding a means of improving memory. Several years ago, he and his colleagues discovered a class of molecules they called am-

pakines. These drugs increase the efficiency of glutamate transmission at AMPA receptors, thereby allowing weaker stimuli to activate NMDA receptors. When hippocampal slices were treated with these drugs, LTP induction was facilitated. And when rats were given the drug, hippocampal-dependent learning speeded up. Michael Rogan and I teamed up with Ursula Staubli to test the effect of these substances in fear conditioning in rats. We found that in the presence of ampakine treatment, the rats learned faster. It thus appears that enhancing glutamate transmission in both the amygdala and hippocampus facilitates learning. Lynch is currently working on ways to apply these potentially powerful agents to human memory in the hope of developing memory-enhancing drugs.

Recently, Joe Tsien of Princeton genetically engineered mice so that their NMDA receptors worked more efficiently. These mice showed faster learning of both a spatial task and fear conditioning. Tsien nicknamed the strain "Doogie" (after the precocious doctor of prime-time fame). This finding provides powerful support for the role of NMDA receptors in learning, and also highlights the possibility that the genetic manipulations might someday make possible the rescue of the loss of memory function in aging humans.

The essence of who you are is stored as synaptic interactions in and between the various systems of your brain. As we learn more about the synaptic mechanisms of memory, we learn more about the neural basis of the self.

CHAPTER SEVEN

THE MENTAL TRILOGY

BETWEEN KNOWLEDGE AND DESIRE STANDS . . . FEELING.

—Immanuel Kant

Throughout much of history, the mind has been viewed as a trilogy; a tripartite amalgam that includes cognition, emotion, and motivation) For some, the trilogy was a description of different aspects of a single mental faculty; whereas for others, it represented three distinct, separate capacities. During most of the twentieth century, both versions of the mental trilogy were out of favor. When the behaviorists reigned, psychology ignored the mind altogether, making the mental trilogy moot. Later, the cognitive revolution brought the mind back to psychology, but thinking and related cognitive processes were (and for the most part still are) emphasized at the expense of emotion and motivation. Clearly, however, it is important to understand not just *how* we attend to, remember, or reason, but also *why* we attend to, remember, or reason about some things rather than others. Thinking cannot be fully comprehended if emotions and motivations are ignored.

In previous chapters, we've seen how neuronal circuits are assembled during development, and how these circuits are modified when we learn and remember. Now we will begin to use this basic information about circuits and their plastic properties to explore broader aspects of mental function, that is, to begin to assemble a neurobiological view of the self. In doing this, we will consider each component of the mental trilogy, as well as interactions among them. Thinking is the subject of this chapter, and emotion and motivation the following two.

MENTAL JUGGLING

An idea, an image, a sensation, a feeling: each is an example of what psychologists call mental content—stuff that is in the mind. Mental content was the subject matter of experimental psychology when it first emerged as a discipline in the late nineteenth century. But John Watson and fellow behaviorists replaced this focus on subjective states with a mind-less psychology of objectively measurable events (stimuli and responses). When the cognitive revolution later made the mind fair game again, it did not do so by reviving subjective psychology. The thinking process itself, rather than the conscious content that results from thinking, became, and largely remains, the subject matter of cognitive science.

"In order for a mind to think, it has to juggle fragments of its mental states." This simple statement by Marvin Minsky, one of the architects of the branch of cognitive science known as artificial intelligence, gets right to the heart

of the matter. Imagine, as Minsky suggests, rearranging the furniture in a room familiar to you. You shift your attention back and forth between locations. Different ideas and images come into focus, and some interrupt others. You compare and contrast alternate arrangements. You may concentrate your entire mind on a small detail one moment, and on the whole room the next. How does the mind do this juggling, and how does it keep track of the imaginary changes? The answer is that the mind uses something called working memory.

How many times have you looked up a phone number and then forgotten it after being momentarily distracted? The reason for this is that you put the number into working memory, a mental workshop that accommodates one task at a time. As soon as a new task engages working memory, the content of the old task is bumped out. For that reason, unless you keep rehearsing the phone number and manage to ignore other things that compete for your attention, it will not remain in your mind.

Working memory is one of the brain's most sophisticated capacities and is involved in all aspects of thinking and problem-solving. It allows you to read a menu and keep its various options in mind while also considering the specials announced by the waiter and then return to the thought you were having before the waiter appeared. It underlies your ability to hold a conversation, play board games like chess, or direct yourself to an unfamiliar destination on the basis of having just looked at a map. In addition to being used in such routine daily activities, working memory also contributes to special human endeavors, like composing music or solving complex mathematical problems, or any other situation in which information has to be held in mind in order to complete a task.

Our understanding of working memory owes much to the pioneering work of Alan Baddeley in the early 1970s. On the basis of his studies of short-term memory, he came to view the mind in terms of two kinds of cognitive systems: a set of specialized systems dedicated to specific mental tasks, and a general-purpose system utilized in all active thinking processes.

Specialized systems come in two flavors. Verbal systems, like systems involved in speech comprehension, are mainly present in the human brain, whereas nonverbal systems are present in all brains. Nonverbal specialized systems are epitomized by sensory systems. Each is involved in processing unique kinds of stimuli (sights, sounds, smells, and so on). As part of their operation, the verbal and nonverbal specialized systems are able to retain what they've just processed for brief amounts of time (seconds). This capacity aids in perception, allowing the system to compare what it is seeing or hearing now to what it saw or heard a moment ago. For example, when listening to a lecture, you have to hold the subject of each sentence in your mind until the verb appears, and sometimes you have to refer back to your memory of earlier sentences to figure out the referent of a pronoun.

The general-purpose system consists of a workspace and a set of mental operations called executive functions that are carried out on information held in the workspace. Although only a limited amount of information can be retained at any one time, the workspace can hold on to and interrelate information of different types from different specialized systems (the way something looks, sounds, and smells can be associated with its location in external space and with its name). This ability to integrate information across systems allows for abstract representation of objects and events. It is especially well-developed in humans, and is likely to contribute to the uniqueness of human cognition.

The information in your working memory is what you are currently thinking about or paying attention to. And because working memory is temporary, its contents have to be constantly updated. But working memory is not a pure product of the here and now. It also depends on what we know and what kinds of experiences we've had in the past. In other words, it depends on long-term memory.

The importance of long-term memory to thought cannot be overemphasized. One of the earliest examples of its significance is still one of the best. In the 1930s, Sir Frederic Bartlett had people listen to folktales from foreign countries and later asked them to recount the stories. Not surprisingly, he found that these unfamiliar stories were not remembered very accurately. What was surprising was that the errors of recall were not random but were quite systematic. The subjects often rewrote the stories in their own minds, especially parts that were particularly foreign to them, revising the plot to the point where it resembled a more familiar Western narrative. To explain his findings,

Bartlett proposed that “Remembering is ... an imaginative reconstruction, or construction, built out of the relation of our attitude towards a whole active mass of past experiences.” He concluded that when we face a problem, we draw upon mental schemata, organized bundles of stored knowledge. For example, if you are asked a question about how baseball is played, you would draw upon a baseball schema, your collective knowledge of baseball obtained from specific direct experiences you’ve had with the game as spectator or player, as well as things you’ve heard or read about baseball. Bartlett’s findings do not just concern the personal, idiosyncratic, and fallible nature of memory, but also emphasize how long-term memories, when retrieved into the temporary workspace of working memory, can guide our thoughts and actions, as well.

It has been known for centuries that we can only keep a few things active in our minds (in working memory) at once.” George Miller, one of the pioneers in cognitive psychology, figured out, through psychological experiments, that the magic number is about seven pieces of information. Some people can hang on to eight or nine, whereas others manage only five, but, on average, temporary storage can hold about seven items. (It’s probably no coincidence that telephone numbers within an area code were designed to have seven digits.) But, as Miller noted, we can effectively expand that capacity by chunking or grouping information—it’s about as easy to remember seven letters as seven words or ideas. No doubt one of the reasons human cognition is so powerful is because we have language in our brains, which exponentially increases the ability to categorize information, to chunk. A whole culture, for instance, can be implied by a name.

The concept of working memory subsumes what used to be called shortterm memory. But as the term *workspace* implies, working memory is more than just an area for temporary storage. It underlies mental work. As Minsky noted, thinking involves juggling of mental items—comparing, contrasting, judging, predicting. It is the job of the executive functions of working memory to do the juggling.

In the spirit of viewing the mind in terms of computer-like operations, some cognitive scientists like Tim Shallice and Phillip Johnson-Laird have referred to executive functions as supervisory or operating system functions. A computer operating system is responsible for controlling the flow of information processing, moving information from permanent memory (ROM) to a central processing unit with active memory (RAM), scheduling tasks to be performed using the active memory, and so on. Similarly, executive functions are involved in the constant updating of temporary memory, selecting which specialized systems to work with (pay attention to) at the moment, and then moving relevant information into the workspace from long-term storage by retrieving specific memories or activating schemata pertinent to the immediate situation. Through executive functions, specialized systems are also directed to attend to certain specific stimuli and to ignore others, depending on what working memory is working on. In complex tasks involving multiple kinds of mental activities, executive functions plan the sequence of mental steps and schedule the participation of the different activities, switching the focus of attention between activities as needed. Executive functions are crucially involved in decision-making, allowing you to choose between different courses of action given what is happening in the present, what you know about such situations, and what you can expect to happen if you do different things in this particular situation. Executive functions, in short, make practical thinking and reasoning possible.

The executive represents a powerful mental capacity, but is not all-powerful. Like the workspace, it has its limits. It basically can do one or at most a few things at a time. This is why you forget a phone number if you are distracted while dialing. With practice and training, we can learn to divide our attention between two mental tasks simultaneously, but only with difficulty. In this sense, the executive is more like an old-fashioned DOS operating system that can only run one program at a time than like a multitasking Windows operating system that can concurrently run word processing, spreadsheet, e-mail, calendar, and other programs.

But there’s also a sort of chunking that takes place in executive functions. As we’ve seen, the executive is involved in scheduling the sequence of steps in a complex task. Here, the executive is doing more than one thing at a time, but the things are all related to the overall goal. If the executive has to work on multiple unrelated goals at the same time, however, the system begins to

fall apart, especially if the goals conflict with one another. An easy way to stress people is to make them do too much at once. Planning, decision-making, and other aspects of mental life suffer when the executive is overloaded.

HEADQUARTERS

Aleksandr Romanovich Luria, the great Russian psychologist, developed broad-ranging and influential theories of brain function on the basis of his studies of World War II soldiers who had sustained gunshot wounds to the head. One of his major conclusions was that damage to the frontal lobe interferes with the ability to plan and execute goal-directed behavior. Other studies of neurological patients with frontal lobe damage, as well as of psychiatric patients who underwent frontal lobotomy, supported the conclusion that the frontal lobes are involved in executive functions (planning, problem-solving, and behavioral control), as well as in short-term or temporary

ory. More recently, it has been shown that neural activity, as measured with devices such as PET and MRI scans, increases in the frontal cortex when humans perform tasks that require temporary storage and executive functions.” Working memory thus has come to be thought of as a function of neural circuits in the frontal lobes.

The frontal lobes (one sits on each side of the brain) are huge, and account for about one-third of the mass of the human brain.” All mammals have frontal cortex, but, for most creatures, its main job is movement control. The prefrontal cortex, which is located in front of the movement-control regions, is especially well developed in primates (some say this region doesn’t even exist in nonprimate mammals). This is the region classically implicated in working memory.

As in most areas of neuroscience, animal research is crucial for determining exactly how neural systems participate in psychological functions, and much of our current understanding of the role of the prefrontal cortex in working memory has come from studies of monkeys. The first clues emerged in the 1930s, when it was found that damage to the monkey prefrontal cortex interfered with the ability to remember under which of two objects a reward (for example, a raisin) was hidden during a delay period in which the objects were out of sight.²³ This temporary storage deficit in the so-called delayed response task in animals with prefrontal damage has been extensively studied and has been repeatedly confirmed in various ways using different kinds of stimuli.

The role of the monkey prefrontal cortex in working memory has been elaborated in work by the laboratories of Joaquin Fuster at UCLA and Pat Goldman—Rakic at Yale. Both researchers have recorded the electrical activity of prefrontal neurons while monkeys perform delayed response tasks and other tests requiring temporary storage and have shown that cells in this region become particularly active during the delay periods. It is likely that these cells are directly involved in retaining information during the delay.

The synaptic connections of the prefrontal cortex, also elucidated through studies of monkeys, explain its ability to participate in working memory tasks involving many different kinds of stimuli) The prefrontal cortex is a convergence zone. It receives connections from various specialized systems (like the visual and auditory sensory systems), enabling it to be aware of what’s going on in the outside world and to integrate the information it gathers. It also receives connections from the hippocampus and other cortical areas involved in long-term explicit memory, allowing it to retrieve stored information (facts, personal experiences, schemata) relevant to the task at hand. In addition, it sends connections to areas involved in movement control (including movement-control areas in the frontal cortex as well as in subcortical regions), allowing executive decisions to be turned into voluntary actions.

A VISION OF WORKING MEMORY

The visual system is one of the most thoroughly studied systems in the brain. For this reason, research on working memory and other cognitive processes often uses visual stimuli. Studies of visual processing, combined with the vast amount of knowledge about the anatomical connections and function of the visual system, have thus helped

construct a fairly detailed understanding of the synaptic connections underlying working memory.

Cortical visual processing begins in the primary visual area located in the occipital lobe (the rearmost part of the cortex). This area receives visual information from the visual thalamus, processes it, and then distributes its outputs to a variety of other cortical regions. The work of Leslie Ungerleider, Semir Zeki, and David Van Essen, among others, has been responsible for elucidating the pathways of visual processing in the cortex. Although these circuits are enormously complex, the pathways responsible for two broad aspects of visual processing are fairly well understood. Thanks to Ungerleider and Mort Mishkin, these have come to be known as the “what” and “where” visual pathways. The “what” pathway is involved in object recognition, and the “where” pathway in figuring out the spatial location of that object relative to other stimuli in the outside world. The “what” pathway involves a processing stream that travels from the primary visual cortex to the temporal cortex, and the “where” pathway goes from the primary cortex to the parietal cortex.

The end stage of the “where” pathway in the parietal cortex is directly connected with the prefrontal cortex. As Goldman-Rakic, Fuster, and others have shown, cells in the prefrontal region are active during delay periods in spatial tasks (which involve “where” processing), indicating that these cells may be involved in the temporary storage of the stimulus information during the time when the stimulus is absent. Moreover, cells in the “where” processing area in the parietal cortex are also active during the delay. Synaptic interactions between the parietal “where” area and prefrontal cortex may therefore allow the animals to keep in mind the spatial location of the reward.

The “what” area in the temporal cortex is also connected with the prefrontal cortex. Earl Miller and Bob Desimone therefore asked whether cells in the “what” visual pathway and the prefrontal cortex might both be active during a delay period in which monkeys had to remember which of two similar stimuli was associated with reward. In this experiment, the reward was paired with a particular stimulus, which is why the “what” stream of processing was involved. They found that cells in both the prefrontal cortex and the temporal cortex “what” processing area were active during the delay.

The maintenance of visual information in working memory thus appears to depend crucially on information transfer over synaptic pathways between specialized areas of the visual cortex and the prefrontal region. The pathways from the specialized visual areas tell the prefrontal cortex “what” is out there and “where” it is located. Moreover, these are two-way streets: the prefrontal cortex, by way of synaptic pathways back to the visual areas, instructs the visual areas to attend to and stay focused on those objects and spatial locations that are being processed in working memory. Recent studies by Goldman-Rakic and Liz Romanski have found that auditory working memory involves similar relations between auditory processing streams and prefrontal areas, suggesting that this scheme involving linkages between specialized sensory processing systems and prefrontal cortex may be generally applicable across many systems.

ORDERS FROM THE TOP

The flow of information from lower- to higher-level processing stations is what cognitive psychologists call bottom-up processing, and the flow from higher to lower stations is called top-down processing. Thus, the pathway from the “what” area to the prefrontal cortex transmits bottom-up information (placing object-identifying information in the workspace), and the pathway from prefrontal cortex back to the “what” area transmits top-down information (executive control signals that help keep attention focused on the object that is represented in the workspace). Top-down activities are just another name for executive functions.

Jonathan Cohen and David Servan-Schreiber used the so-called Stroop task to study top-down executive influences on lower-order processing. In this procedure people are given brief visual presentations of simple words (color names) and are asked to either name the word or state the color of ink the word is printed in. On some trials, the word name and color are in conflict (the word *green* is presented in red letters) and on others they are congruent

(the word *green* is presented in green letters). It takes much longer to state the ink color in the conflict than in the congruent situation, but subjects tend to get it right. Unlike normal subjects, patients with damage to the prefrontal cortex fail this task). They are unable to use the instruction given by the experimenter to suppress the more commonly given response of wordnaming. Schizophrenics, who are believed to have frontal dysfunction, also do poorly. Functional imaging studies by several groups have shown that the prefrontal cortex is activated during the conflict condition in normal subjects but not in schizophrenics. These findings suggest that executive functions, by way of connections from the prefrontal cortex back to lower-order processors (color and word processors), use the instructions given by the experimenter (and kept active in working memory) to inhibit the more natural word-naming response and facilitate the less common response of stating the ink-color name. Cognitive scientists often find it useful to program computer simulations to help them understand how psychological and sometimes neural processes work. Cohen and Servan-Schreiber did this in an effort to account for the role of the prefrontal cortex in the Stroop task.

Although it generally has been assumed that temporary storage and executive functions are carried out by the same networks in the prefrontal cortex, most studies have focused on the former rather than the latter. However, Mark D'Esposito and colleagues designed an experiment to explore explicitly the role of prefrontal areas in executive functions. They had human subjects perform a verbal identification task (e.g., respond if you hear a vegetable name) and a visual task (respond if a certain spatial arrangement of stimuli occurs).

The tasks were performed either separately (which was easy and did not require executive functions) or simultaneously (which was difficult and did require executive functions, since the subjects had to keep track of two different kinds of stimuli at once). The results showed that the prefrontal cortex was activated when the two tasks were performed simultaneously, but not when the tasks were performed individually. Since there was no temporary storage requirement in the task (no delay between stimulus and response), the results specifically implicate the prefrontal region in executive functions. A number of other studies have also now clearly demonstrated that the human prefrontal cortex is indeed activated when executive operations are performed, as well as when information is stored temporarily.

For obviously reasons, it is more difficult to study executive functions in nonhuman primates. You can't just give verbal instructions to a monkey, which has to be trained nonverbally on each step of a given task. But it has long been thought that deficit in delayed response performance by monkeys with frontal lesions was not simply due to a loss of temporary storage, since the deficit was worse when there were distracting cues. It was as if the monkeys couldn't ignore irrelevant, distracting information.

Recently, studies have begun to provide evidence for the operation of executive-like functions in the monkey brain. For example, John Reynolds and Bob Desimone trained the monkeys to focus their attention on the location on a screen where a dot appears. This is, in essence, the training of an executive function, attention. Then they presented a picture in that same area (that is, within the field of attention). Cells in the early stages of the visual cortex responded to the picture. When a second stimulus was presented on the screen with the first, though, the cells did not respond. Executive functions continued to focus cellular activity on the first picture, and the new information about the second picture was ignored. This work, like the studies described above, suggests that attentional signals are involved in directing traffic in the "what" pathway, controlling which stimuli are processed. In related work, Leslie Ungerleider and colleagues had human subjects perform similar tasks while functional brain activity was measured. They reached a very similar conclusion: that top-down activity regulates lower visual areas.

Studies by Earl Miller's lab suggest that the prefrontal cortex in the monkey brain, as in the human brain, is the source of top-down signals. Monkeys were first shown a target stimulus. That stimulus then appeared in an array with two irrelevant stimuli. After a delay, the three-stimulus array appeared again. If the target stimulus was in the same location, the monkey had to make a response on a lever. The results showed that prefrontal cells were particularly responsive to the selection of the target stimulus from the array.

Further, although cells in the visual cortex are also sensitive to target selection, their activity occurred after prefrontal activity, suggesting that the attentional signal originates in the prefrontal region and is transmitted in top-

down fashion to the visual cortex.

But is the activity occurring in the prefrontal cortex necessary in cognition, or just along for the ride? In other words, is the frontal cortex activity the result or the cause of top-down cognition? Results obtained by Robert Knight suggest that it is necessary, that it is causal. Knight recorded electrical activity in the visual cortex of humans using electrodes placed on the scalp. Although this technique does not provide detailed information about specific neural circuits, it is sufficient to detect activity in areas like the visual cortex. Normally, if subjects are instructed to expect a certain stimulus, the electrical response to that stimulus in the visual cortex will be enhanced. However, this attentional modulation of lower-level sensory processing did not occur in patients with damage in the prefrontal region. The prefrontal cortex appears to be necessary for top-down cognition to work.

The importance of top-down signals from the prefrontal cortex in selecting neural activity in lower processors was also demonstrated in an elegant study by Yasushi Miyashita's team in Japan. They made lesions in the brains of monkeys in such a way that allowed them to present visual stimuli so that the "where" area in the temporal cortex received information either bottom-up (from the lower areas of the visual system) or top-down (from the prefrontal cortex), but not both. They found that top-down information from the prefrontal area was sufficient to selectively activate cells in the "where" pathway.

A central aspect of executive functions is decision-making. Studies by several research groups have begun to reveal how this works in the brain. Most of the investigation has focused on a simple behavior, eye movement, which is one that offers many advantages to the experimenter. A monkey is first instructed, through behavioral training, to respond to certain stimulus conditions by moving its eyes to a certain location. This response is then used to assess how the monkey makes decisions about what to respond to on the basis of what it sees. For example, Michael Shadlen and Jong-Nam Kim showed monkeys visual stimuli moving in one direction or another. The moving stimuli were embedded in a busy background, making it impossible to respond correctly every time. The monkey, in other words, had to use the imperfect information available to make a decision about the stimulus. Further, a delay was imposed between the stimulus and the response. Earlier work had shown that the motion stimulus is processed in a special area of the visual cortex called MT, which has connections with the prefrontal cortex. In the new experiment, Kim and Shadlen recorded neural responses in the prefrontal cortex and found that neural activity in this region predicted, during the delay, which direction the eyes were about to move in. They concluded that these prefrontal cells are involved in the conversion of stimulus information into response plans. These cells are, in their view, part of the decision-making process.

Some investigators argue that circuits in the so-called "where" area of the parietal cortex, like prefrontal cells, are involved in planning and decision-making about movements. In many ways, these cells apparently function similarly to prefrontal cortex cells. However, one difference seems to be that the parietal circuits are involved in planning the next movement, whereas prefrontal circuits participate in planning several steps ahead. Later, when we consider emotion and motivation, we will have the opportunity to explore other aspects of decision-making.

ONE SYSTEM OR MANY?

Studies by Joaquin Fuster first suggested that the prefrontal cortex might integrate spatial ("where") and object ("what") information, since cells responding to these two kinds of information were intermixed in the prefrontal cortex. This result was consistent with the view that the prefrontal cortex engages in general-purpose temporary storage across many processing domains. More recently, though, Pat Goldman-Rakic has questioned the assumption that the prefrontal cortex is a general-purpose working memory processor. She and Fraser Wilson found evidence that cells involved in object and spatial processing are clustered in different areas of the prefrontal cortex. Recently, Liz Romanski and Goldman-Rakic found additional areas that they believe are involved in auditory working memory. These findings suggest that temporary storage is carried out by domain-specific regions in the prefrontal cortex and question the notion of a universal workspace. Ed Smith and John Jonides reached a similar conclusion, on the basis of a survey of functional imaging studies of working memory in humans. However, Earl Miller has pointed out that

there are a number of cells in each frontal area that do not respond exclusively to one kind of stimulus. Further, the different areas are interconnected extensively with one another. It is quite possible that the various areas work together to integrate information across domains, and could, in fact, constitute a single distributed system (a system that is spread over several brain regions) that mediates temporary storage, rather than a collection of independent systems. Additional work is needed to resolve this issue.

Regardless of how temporary storage functions are laid out, it appears that executive functions are not partitioned in the prefrontal cortex on the basis of the stimulus domain. This does not mean, however, that there is a single executive area that takes care of all executive functions. Executive functions, in fact, seem to be spread across multiple regions of the frontal cortex. For example, the classic area of the prefrontal cortex implicated in working memory has been the lateral prefrontal cortex, located on the outside surface of the frontal cortex. However, another area, located on the medial or inside surface of the frontal cortex, has also been consistently activated in functional imaging studies of working memory, and is one of the main regions implicated in executive functions. This region, the anterior cingulate cortex, like the classic working memory area in the lateral prefrontal region, receives inputs from various specialized sensory systems, and is anatomically interconnected with the classic lateral area. Moreover, both regions are part of what cognitive psychologist Michael Posner has called the frontal lobe attentional network, a cognitive system involved in selective attention, mental resource allocation, decision-making processes, voluntary movement control, and/or resolving conflict between competing stimuli. It is tempting to think of the executive aspects of working memory as involving synaptic interactions between neurons in the lateral prefrontal and anterior cingulate regions, and perhaps other areas. Nevertheless, the various aspects of executive functions may not be equally distributed across the various areas. There is, in fact, some evidence for localization of particular executive functions—that is, in some tasks different aspects of executive functions (for example, stimulus or response selection vs. conflict resolution vs. decision-making) have been found to engage different areas of the prefrontal cortex to different degrees. This in no way implies the existence of multiple executives, each with their own planning and decision-making capacities. Instead, it suggests that the various component executive functions are achieved by a set of interconnected circuits that are spread over several brain regions in the frontal cortex, and even other regions, as discussed more in chapter 9.

THE NUTS AND BOLTS OF THOUGHT

How does working memory function at the level of cells and synapses? While this process is still not fully understood, we can at least begin to piece together an explanation.

The prefrontal cortex, like other areas of the neocortex, has six layers. And, as in other areas, the middle layers tend to receive inputs from other regions, while the deep layers tend to send outputs to other regions. So axons from other cortical areas, such as areas involved in “what” and “where” processing, form synapses on cells in the middle layers of prefrontal cortex. These input cells then send axons to cells in the deep layers, which give rise to connections that go back to the middle-layer cells, or to other cortical or subcortical areas, especially areas involved in the control of movement, and thus of behavioral responses. In this manner, the deep and middle layers can influence each other. In addition, though, input cells in the middle layers and output cells in the deep layers each give rise to local connections to other cells in the same layer. This arrangement allows the input cells to influence other input cells, and output cells to influence other output cells.

Transmission of inputs to and outputs from the prefrontal cortex, and between cells within and between layers within the prefrontal cortex, is mediated by the binding of presynaptically released glutamate to postsynaptic receptors. Interestingly, extrinsic inputs to these circuits account for only a small part of the excitatory synaptic connectivity of the prefrontal cortex. The connections within the prefrontal cortex, both within and between layers, are far more numerous than the connections coming in from other areas, such as sensory processing regions. The mutual excitations mediated by the internal connections enable input signals from the outside to be amplified and kept active, and may well contribute to the sustained activity that has been observed during delay periods:

Of course, there is also inhibition in these circuits. Most of the inhibitory connections are on the deep-layer output cells. However, because these cells also have excitatory projections back to upper layers and to other cells in the same layer, local inhibition in the lower layers can regulate much of the excitatory flow through the network.

Much of what we know about the contribution of excitatory and inhibitory circuits in the prefrontal cortex to working-memory centers around the neuromodulator dopamine. The prefrontal cortex receives a rich supply of axons containing dopamine, and depletion of the substance in the prefrontal cortex (by injecting dopamine-destroying drugs) is as effective as complete removal of the prefrontal cortex in disrupting delayed-response performance in monkeys. In dopamine-depleted animals, performance recovers if a drug that activates a certain class of dopamine receptors is infused into the prefrontal cortex, thus tricking the receptors and the postsynaptic cells into acting as if dopamine fibers were still there releasing their transmitters. Infusion of dopamine into the prefrontal cortex of young monkeys enhances working memory capacity, and in older ones, it reverses the age-related decrement in working memory.

Dopamine cell bodies are located in the brain stem, in a region called the ventral tegmental area. The axons of these cells then branch extensively and reach many areas of the forebrain, including the prefrontal cortex, where their terminals release dopamine. In primates, the dopamine terminals are fairly evenly distributed throughout the layers, allowing dopamine to bind to receptors and then modulate excitatory and inhibitory transmission in both the input and output layers. Although there are many subclasses of dopamine receptors, the D_1 family (which includes D_1 and D_5 receptors) has been most clearly implicated in working memory. These receptors are located on the spines and shafts of dendrites of excitatory cells and seem to reduce the transfer of excitation from the dendrites to the cell bodies, allowing only especially strong excitatory inputs to get through to the cell body and elicit excitation. Dopamine release in the prefrontal cortex also seems to facilitate GABA inhibition, possibly by way of presynaptic facilitation of transmitter release, leading to a further reduction of excitation through prefrontal circuits. Some of these effects appear to involve the triggering of protein kinase A in cells containing dopamine receptors. Integrating these findings, Amy Arnsten has proposed that dopamine participates in working memory by biasing cells to mainly respond to strong inputs and thereby focusing attention on active current goals and away from distracting stimuli.

THE WIZARD OF OZ

Most of us have an intuitive sense that thinking (cognition) and consciousness are closely related. This intuition is correct but incomplete. Virginia Woolf was probably referring to this when she pointed out, "We all indulge in the strange, pleasant process called thinking, but when it comes to saying..., what we think, then how little we are able to convey! Imagine that you are mentally engaged in some activity, like reading the newspaper. You ignore the other things going on around you. But if something meaningful occurs in the background, like the sound of your name, you stop reading and turn to the person who called you. Although your conscious mind was ignoring everything apart from the visual signals from the paper, your brain was not. Inputs from other sensory systems continued to be actively processed, otherwise the mention of your name could not have interrupted you. Cognition and consciousness, therefore, are not the same. What, then, is consciousness, and what determines what we are conscious of?

Consciousness can be thought of as the product of underlying cognitive processes. We've actually already seen how the key processes work. The stuff we are conscious of is the stuff that working memory is working on. Though cognitive science provided a way of studying the mind without getting entangled on the controversial question of consciousness, it has, in the process of accounting for working memory, also provided a practical approach for understanding how consciousness works.

Let's translate the example above into cognitive language. During focused attention, mental resources are allocated to the task at hand. The executive, in other words, keeps lower-level processors engaged in activities that support the task being worked on. However, if lower-level processors that are not being worked with at the moment detect some event that is unrelated to the current task but is more important than the current task, resources are allo-

cated to processing the new event. The task management, scheduling, and conflict resolution functions of the executive shift attention to the new event and move information relevant to it into working memory. This forces out the existing information. Behaviors relevant to the new task are initiated: you turn toward the source of the sound. If, after saying your name, the person asked, "What's that smell?" you would probably start sniffing in an effort to answer the question. During these operations, memories are retrieved into working memory, enabling you to recognize that the sound spoken was your name and that it was your friend who spoke it. (How else would you know your name except by virtue of the fact that you remember it; and how else do you know that it was your friend who said your name except that you remember her voice and appearance?) Retrieved episodic and semantic memories, as well as schemata, then aid in executive decision-making.

Although the end result of executive processes (monitoring, resource allocation, task management, conflict resolution, memory retrieval, etc.) was the representation of your name as conscious content in working memory, it's important to recognize that the executive processes that made this possible functioned unconsciously. As neuroscience pioneer Karl Lashley pointed out in the early 1950s, we are never aware of processing, but only of the consequences of processing. Like the Wizard of Oz, executive processes work behind the scenes.

The working-memory model of consciousness suggests that the prefrontal cortex should play some role in consciousness. This is not a novel idea. For example, the Russian psychologist Luria, whom we encountered earlier, pointed out two characteristic deficits of patients with damage to the frontal lobes: they lack spontaneous, purposeful behavior and they lack an understanding of their deficits. And Arthur Benton, an American neurologist who long studied the effects of frontal lobe damage, described a typical patient in the following way: "He has never been seen to sit self-absorbed, to daydream or to indulge in introspection. These patients are not, strictly speaking, "unconscious," but rather seem to have a fractured understanding of what they. At the same time, they do not typically suffer total destruction of the prefrontal cortex, and the remnants of their conscious states may be due to the undamaged parts of this large and complex brain region—as we've seen, working-memory functions are not located in a single region of the prefrontal cortex but are distributed across widespread regions.

On the other hand, there may exist primitive levels of consciousness, especially involving the passive awareness of events as opposed to the active use of on-line information to guide decision-making and behavior, that do not depend on the prefrontal cortex. These kinds of mental states may typify consciousness, to the extent that it does exist, in organisms that have less or no prefrontal cortex, as discussed below.

If the prefrontal cortex plays an essential role in human consciousness, it should be involved in explicit memory, memory to which we have conscious access. Indeed, damage to the human prefrontal cortex disrupts the conscious retrieval of long-term memories, especially episodic memories. The prefrontal cortex is activated during episodic memory retrieval, and there is also evidence that the prefrontal cortex is involved in the encoding (formation) of episodic memory. These findings suggest that, in order to have an explicit, conscious memory, in addition to the involvement of the medial temporal lobe (chap. 5), two conditions need to be satisfied: you have to be conscious of the information constituting it at the time of the original experience (that is, the experience had to be represented in working memory at the time it occurred), and during retrieval, you have to transfer the information from cortical storage circuits (chap. 5) into working memory.

Just as researchers studying working memory often use visual stimuli, theorists interested in the brain and consciousness tend to focus on visual awareness as a window on conscious experience. For example, Nobel laureate Francis Crick and his colleague Christof Koch have proposed an explanation of visual awareness that builds upon knowledge of the role of the visual system in working memory. They argue that we have access to information processing occurring in visual-processing areas connected with the working memory areas of the prefrontal cortex but do not have access to information processed by areas not connected to these areas. Indeed, the primary area of the visual cortex is not connected with the prefrontal cortex, and we do not experience the contours and shadings of stimuli that it processes. However, the later processing areas are connected with the prefrontal cortex and we experience the products of these areas—global features of objects, like their shape, color, motion, and location.

Several studies provide support for the Crick/Koch theory. For example, Roger Tootell and colleagues examined whether the brain networks involved in the processing of visual motion would be active when subjects were tricked into seeing motion where none exists. They used the waterfall illusion, which is produced by the presentation of stationary stimuli at slightly different times. The result was indeed that the motion-processing area, which is an advanced visual area that is connected with the prefrontal cortex, was particularly active during the illusion of movement.

Given that different features of a visual stimulus (shape, color, location, motion) are processed in different cortical areas, in order to have a conscious perception of the stimulus as a whole object, rather than as a collection of features, the elements have to be bound together. Although we can separately identify these features, we usually experience them together as the coherent object rather than as component parts. Understanding how the integration occurs is called the binding problem).

Binding is taken care of, at least in part, by the transfer of information from the individual processing regions to areas that integrate the information. As we've seen, this kind of convergence occurs within each sensory system. For example, the earliest stages of the visual cortex process primitive features of objects (contours and brightness, for example) and the later stages process more complex features ("what," "where," and movement, for example). Further, different complex features are processed in different circuits, requiring integration or binding across circuits as well as within them. And visual stimuli do not occur in isolation but in the context of other kinds of stimuli (sounds and smells, for example), requiring binding across sensory modalities. Further, the conscious experience of a sensory stimulus is not just a sensory experience. We usually experience stimuli as meaningful objects rather than as raw sensations. Thus, the way the stimulus looks, sounds, and smells has to be integrated with relevant information stored in memory, including facts and past experiences, as well as stored information about the emotional and motivational significance of the stimulus. Because prefrontal regions receive convergent inputs from sensory, memory, emotional, and motivational circuits, they are believed to be capable of performing the complex kind of information integration (binding) that must occur in the brain during a conscious experience.

But some researchers feel that information convergence in working memory is not sufficient to account for the conscious experience of a stimulus, or for our ability to make decisions and take action relevant to that stimulus. They argue that the additional ingredient needed is neuronal synchrony. This term refers to the simultaneous firing of populations of neurons. Synchrony is proposed to achieve two important goals—enhanced activation of postsynaptic cells and coordination within local areas across widespread regions. It is well-established that a group of postsynaptic cells are more strongly activated when they receive synchronous inputs from presynaptic cells. Thus, if the brain is processing a salient visual stimulus, cells in the visual areas will fire synchronously. As a result, the information processed in these areas will activate prefrontal areas more strongly, and the visual stimulus will make it into working memory more easily. The other proposed role of synchrony is coordination. That is, synchrony is believed by some to help solve those aspects of the binding problem that information convergence cannot. According to this view, neurons that fire together in widespread brain regions are temporarily bound together, and this coherence of firing, when combined in just the right way across the brain, facilitates the representation in working memory of momentarily relevant information from diverse regions. While considerable evidence exists showing that synchrony occurs and enhances postsynaptic firing, the second role of synchrony, coordination across widespread regions, is controversial, especially as an explanation of conscious perception. (We'll encounter a less controversial role of synchrony in chapter *ii*, namely in the coordination of synaptic plasticity across brain regions.)

The working memory theory of consciousness, even without the assistance of synchronous information coordination across brain regions, gives us a neural handle on consciousness. But how far can we go with this pursuit of con-

sciousness and the brain without butting heads with the mind-body problem (chap. 2)? After all, it's one thing to say that the information in working memory is what we consciously experience, and another to explain how those experiences emerge from working memory. Similarly, it's one thing to say that voluntary behavior results when

conscious content in working memory initiates executive functions that then lead to muscle movements, and quite another to explain how conscious content is translated into executive control. To a dualist, one who believes in a separate mind substance that has to somehow be merged with physical matter, these are indeed different phenomena. But if, instead, you, like me, believe that the things we describe in mental terms are, in fact, processes going on in the brain, a working memory point of view is an excellent way of attacking the problem. While it may not at this point fully account for all the subtleties of how consciousness and other aspects of the mind arise from the brain, working memory is clearly a concept we can work with well.

THE REVOLUTION OF THE THINKING BRAIN

As I discussed in chapter 3, the prefrontal cortex is especially well-developed in humans, is present in other primates, rudimentary in nonprimate mammals, and doesn't even exist in other creatures. What are the implications of this for understanding working memory and consciousness?

Let's focus first on the difference between nonhuman primates and other mammals. To do so, we have to consider the organization of the prefrontal cortex in more detail—specifically, the distinction between different prefrontal areas. We've already discussed the lateral prefrontal cortex, which also has divisions, as above. The medial prefrontal cortex is located on the inside wall of the hemisphere. Imagine the brain's hemispheres as the two parts of a hot dog bun. The brown part on the outside is like the lateral cortex and the white part on the inside is like the medial cortex. One region of the medial prefrontal cortex is the anterior cingulate, which, as we've seen, is also involved in executive functions. In addition, the ventral prefrontal cortex, especially the orbital cortex, appears to play a role in working memory, particularly for emotional information, as we'll discuss in chapters 8 and 9. While other mammals have medial and ventral prefrontal cortices, primates alone appear to have lateral prefrontal cortex. Thus, one of the major regions involved in working memory in primates clearly does not exist in other animals. The fact that the cognitive capacities of these creatures do not compare with those of primates suggests that the unique features of primate cognition came with the development of the lateral prefrontal region and its integration with existing networks involving the medial and ventral areas. For example, rats can engage in temporary storage, especially about emotionally relevant information, and can focus attention on specific stimuli and so forth, but are far more limited than primates in their capacity to categorize the world, to discriminate among different stimuli and events, to relate or associate things with one another on-line, and to use the results of these cognitive analyses to guide problem-solving and decision-making.

It's important to point out that temporary storage can be carried out in domain-specific systems, like sensory or emotional systems, which accounts for short-term memory in such animals as birds and reptiles. It also accounts for domain-specific short-term memory in mammals. These domain-specific kinds of short-term memory processes may underlie primitive conscious representations in animals that lack elaborate working memory functions. Domain-specific temporary storage may allow, for example, an awareness of significant stimuli, like the sight of a predator, the pain of being injured, the taste of food, or the joy of sex. When the arousal elicited by the temporarily stored information is sufficiently intense, the activity of other systems may be inhibited, allowing the active system to dominate brain function until the arousing situation subsides. It is known, for example, that sexual impulses are inhibited in threatening situations. This kind of arousal could be achieved by way of interactions between, for example, a predator detection system and brain stem arousal circuits (as described in the next chapter). Indeed, Gary Aston-Jones of the University of Pennsylvania and others have shown that brain stem arousal systems underlie vigilance, sustained behavioral engagement in a task. Such connections could make possible focused attention even in creatures lacking a prefrontal cortex and its multimodal integrative capacity and executive functions.

It is, of course, difficult to know if other animals experience anything, since we can't ask them outright, but if they do, it may be something akin to domain-specific sensory consciousness. The key element that distinguishes working memory from sensory consciousness is that the former allows for the simultaneous interrelation of temporarily stored information across domains and the flexible use of such information in decision-making, capacities that

prefrontal circuits seem to make possible. By this analysis, something akin to human consciousness would be present in other animals with well-developed working memory systems (nonhuman primates) but not in other creatures.

But even nonhuman primates lack the unique features of human consciousness. As the novelist and naturalist Annie Dillard says, "It is ironic that the one thing that all religions recognize as separating us from our creator—our very self-consciousness—is also the one thing that divides us from our fellow creatures. It was a bitter birthday present from evolution.

Other than size, which in this case is significant, the human prefrontal cortex has another important advantage over the prefrontal cortex of nonhuman primates: it has access to a processing module specialized for language use. Increased size adds power, but the presence of language does more than just soup up the cortex. (I'm referring here to the kind of grammatical natural language that characterizes essentially every human brain, rather than the other communicative capacities that can occur in other kinds of animals, like chimps and even parrots.) Language radically alters the brain's ability to compare, contrast, discriminate, and associate on-line, in real time, and to use such information to guide thinking and problem-solving. The difference between having only nonverbal working memory and having both verbal and nonverbal working memory is enormous for how the cognitive system works.

It is, in my opinion, the structuring of cognition around language that confers on the human brain its unique qualities. Other animals may be consciously aware, in some sense, of events going on in their world. They may have domain-specific consciousness, or in the case of nonhuman primates, domain-independent nonverbal consciousness, but lacking language and its cognitive manifestations, they are unlikely to be able to represent complex, abstract concepts (like "me" or mine or "ours"), to relate external events to these abstractions, and to use these representations to guide decision-making and control behavior. Some advanced primates do have the ability to visually recognize themselves in a mirror which suggests a sense of self-recognition in the absence of natural language. Interestingly, recent studies in dolphins and whales show something similar, suggesting parallel evolution of this capacity—but questions have also been raised about the interpretation of these findings. In any event, self-recognition does not necessarily imply selfawareness. We know from many kinds of experiments that stimulus recognition is not the same as being consciously aware of the stimulus. Stimulus recognition requires only that an immediately present stimulus match some representation of a similar stimulus in memory. There is no need to have conscious awareness to have recognition.

When I use the term *consciousness*, I usually am referring to the special qualities of human consciousness, especially those made possible by language. (There may be other ways that the human brain differs as well, but language is a particularly important difference.) This does not imply that we are conscious in English or Chinese or Swahili, nor does it suggest that persons who are deaf and mute should be considered cognitively impaired. My point, rather, is about the way the human brain is wired. That is, the emergence of the cognitive capacities underlying language changed the way the brain works, making it possible for human brains to think and experience events in ways that other brains cannot. The addition of language into the human brain involved a revolution rather than an evolution of function.

P.S.

My idea that language embellishes working memory, and in so doing makes human consciousness unique, bears some relation to Mike Gazzaniga's notion of an interpretive system in the left hemisphere that gives rise to the unique properties of human consciousness. In a nutshell, the essence of the interpreter theory is that our conscious awareness of who we are depends on our linguistic interpretation (labeling, categorizing, explaining) of our experiences as we go through life (this is related to the notion of a narrative or constructed conscious self discussed in chapter 2). The similarity of my ideas about consciousness to Gazzaniga's is perhaps not surprising, since he was my

Ph.D. mentor in the 1970s. Mike had been pursuing consciousness in the brain since his own days in graduate school in the 1960s. As part of my Ph.D. research, we did studies of a split-brain patient, known as P. S., that greatly influenced us both. Mike emerged with the theory of the interpreter, and I with a yearning to understand the emotional brain. But we've come full circle. Emotions now play an important role in his theories of consciousness, and, as we see in the next chapter, consciousness, in the form of working memory, has come to be an important part of the way I think about emotions, especially feelings.

CHAPTER EIGHT

THE EMOTIONAL BRAIN

REVISITED

THE ADVANTAGE OF THE EMOTIONS THAT THEY LEAD US ASTRAY

—Oscar Wilde

Emotions can, as Wilde says, lead us astray. But their inclusion in the mental trilogy by great thinkers through the ages comes less from their distracting qualities than from their contribution to defining who we are. A purely cognitive view of the mind, one that overlooks the role of emotions, simply won't do. As we begin to assemble the self in terms of synapses, emotions must play a major role.

THE EMOTIONAL BRAIN AND ITS VICISSITUDES

As soon as pioneering brain researchers of the late nineteenth century identified regions of the neocortex involved in sensory perception and movement control (chap. 3), William James asked whether emotions might also be explained in terms of these particular functions, or whether they were rather the business of a separate, yet undiscovered, brain system.⁷ Being a pragmatist, he proposed a theory of emotion based solely on functions of

sensory and motor mechanisms. Specifically, he argued that emotionally arousing stimuli are perceived by sensory cortex, which activates motor cortex to produce appropriate bodily responses. Emotional feelings then result when our sensory cortex perceives the sensations that accompany bodily responses. Since different emotions involve specific bodily responses, they have distinct sensory signatures and thus feel different. The essence of James's argument is captured by his conclusion that we do not run from a bear because we feel afraid, but instead we feel afraid because we run.

James's theory was quickly refuted by research showing that complete removal of the neocortex failed to disrupt the expression of emotional responses elicited by sensory stimuli, thus suggesting that sensory and motor cortex could not be the key.' Subsequent studies then implicated specific brain areas, many involving subcortical or old cortical regions, as opposed to neocortex, a trend that led to the view that the brain does indeed have a special emotion system. This trend ultimately culminated around mid-century in the famous and still popular limbic system theory of emotion, which inspired many subsequent studies (we'll consider this theory in detail below). By the mid-1960s, though, research on the neural basis of emotion had all but come to a halt, or at least a slow crawl—after decades of concentrated attention, neuroscientists forsook the topic.'

Like psychologists, brain researchers were strongly influenced by the emergence of cognitive science, and emotions were not part of the cognitive game plan. Emotions seemed more a matter of mental content than of mental processing, and were not pursued by those interested in the thinking process. Although cognitive scientists didn't deny that emotions were important, they believed the subject was not relevant to their field. As interest in cognition rose, research on emotion declined in neuroscience.

In recent years, however, emotion, part two of the mental trilogy, has again become a popular topic for investigation in neuroscience. And enthusiasm by brain researchers has in turn helped revitalize interest in emotions by psychologists. Although I discussed some aspects of this new wave of research in *The Emotional Brain*, much has happened since, especially in terms of how animal studies of emotion relate to the human brain. Before we turn to research on the neural basis of emotion, I therefore want to spend some time discussing the relation of animal and human studies of emotion.

THE CREDIBILITY PROBLEM

Emotion researchers, whether in psychology or brain science, have typically sought to account for what most people think of as the essence of an emotion, the subjective experience that occurs during an emotional state—the feeling of fear when in danger, of anger when mad, or of joy when something good happens. This was clearly the goal of emotional brain theories from James through the limbic system concept. However, most of what we know about the detailed brain mechanisms of emotion comes from studies of emotional behavior rather than from studies of feelings themselves. The explanation for this situation is simple. Feelings can be studied in humans, but, for reasons that will be discussed below, they're more difficult to examine in animals. Since, for both practical and ethical reasons, most brain research is conducted in animals, we end up with a gap between what emotion theories are about (feelings) and what brain researchers actually measure (behavior). This gap, in turn, creates a credibility problem for brain research on emotions.

The credibility issue could be overcome if there were some way to investigate feelings in animals. In humans, the main method by which feelings are studied is that of self-evaluation, often in the form of verbal self-report. However, even self-evaluation is problematic as a scientific tool for studying emotions: As Donald Hebb, father of the Hebb synapse (chap. 6), pointed out, outside observers are often more accurate in characterizing emotional feelings than the experiencing subject. Jealousy, according to Hebb, is often readily detected by impartial observers but denied by the jealous person, who may instead characterize his state as indignation or annoyance. Hebb noted that when observers agree and the subject disagrees about an emotional state, the conclusion of the observers is often a more reliable predictor of future behavior.

The traditional way of using self-report to study feelings is to ask people to reflect back on some previous

emotional experience. But assessments from memory are particularly problematic. Daniel Kahneman and colleagues have shown, in a variety of ways, that what we remember about an emotional experience is an imperfect reflection of what was actually experienced.” For one thing, people tend to remember how they felt at the end of an emotional episode rather than how they felt about the whole episode. In one study, patients undergoing a painful medical procedure, colonoscopy, were asked to rate their pain every sixty seconds. For some subjects, the procedure was prolonged for one minute, but under more comfortable conditions (the scope was stationary during this time). Later, after it was completed, they were asked to evaluate the entire experience. For subjects who had the stationary scope for the extra minute, and thus experienced less pain at the end, the remembered experience was less aversive, even though the procedure actually lasted longer. On the basis of this and many other studies, Kahneman concluded that the remembered emotional significance (what he calls the re

membered utility) of an experience doesn’t necessarily reflect the overall experience (its total utility).

Distortions in memory also extend beyond the intensity of experience and include the content of what is remembered. For example, studies by Elizabeth Loftus and others have shown that memories of emotional experiences are often significantly different from what actually happened during them. She reports many instances in which vivid memories of crime scenes turn out to be inaccurate, if unintentionally so; sometimes, again unintentionally, they are completely fabricated. As Bartlett demonstrated long ago, memories are constructions assembled at the time of retrieval, and the information stored during the initial experience is only one of the items used in the construction; other contributors include information already stored in the brain, as well as things the person hears or sees and then stores after the experience.

Because remembered experience is a distortion of actual experience, and because it’s difficult to elicit real emotional experience in laboratory studies, some contemporary researchers emphasize the importance of obtaining immediate or on-line evaluations during real emotional episodes (measures of instant utility), especially in natural settings. Though such studies circumvent the errors of remembered utility and the artificiality of laboratory research, they are extremely hard to implement, as well as being time-consuming, and in the end are still subject to the biases and measurement problems inherent in any method that relies on one’s introspective evaluation of his or her own mental states. The fallibility and subjectivity of introspection are, after all, what triggered the behaviorist revolution in psychology in the early twentieth century. And recall that the cognitive revolution only succeeded in bringing the mind back to psychology because it figured out how to study the mind without relying on introspection—by focusing on mental processes rather than mental content.

But if you want to assess the introspective content of subjective experience, there just aren’t many alternatives to self-evaluation and verbal report. A subjective experience is by definition one that is known directly only by the experiencing person, and verbal descriptions of subjective states are the most direct way of assessing them. While the use of checklists or scales, which allow subjects to indicate their feelings by selecting from choices, sidesteps the need for verbal behavior as the final response, in the end they, too, depend on introspective evaluation of mental states, and inescapably involve the use of words to classify and categorize mental content. For example, the easiest way to identify which emotion you are feeling is to label it verbally as fear, anger, love, or disgust. Some even insist that you don’t really know what you are feeling until you have labeled it.

FEELINGS IN ANIMALS

Some have argued that an animal’s behavior can be used as a kind of nonverbal self-report to assess its feelings. The neuroscientist Jaak Panksepp takes this approach in *Affective Neuroscience*, as does, more controversially, the psychoanalyst Jeffrey Masson in *When Elephants Weep*.⁷ Their logic is as follows:

Since animals and humans behave similarly when emotionally aroused (for example, similar fear responses are expressed by rats and people in the presence of danger), they must experience the same subjective states as well. If so, it would be possible to use behavioral responses in animals as indicators of feelings.

The flaw in this approach, which takes us right back to the credibility problem, is revealed by the results of a study of human heroin addicts. The subjects were allowed to press a button to administer either saline or a high or low dose of morphine through an intravenous tube. The subjects did not know what was in the tube at any given point. Periodically, they were asked to rate how they felt. When a high dose of morphine was in the tube, they pressed vigorously and also reported feeling high. When saline was in the tube, the subjects pressed little and reported feeling nothing. But when the dose of morphine was weak, the subjects vigorously pressed the button in spite of the fact that they reported feeling nothing. Clearly, one would be misled by using behavior as a measure of what was felt in this case, since the subjects behaved but didn't feel. Emotional responses are not always external mirrors of internal feelings, but are rather controlled by more fundamental processes.

The problem is only compounded when we examine the relation of behavior to feelings in different species. Just because two creatures act the same does not mean they have the same experiences when they perform those actions. A beetle that finds itself under the approaching footstep of a human does what a human would do—it tries to escape before the foot lands. A robot can be programmed to do what a human would do if an object is sent flying toward his or her head—raise an arm to deflect the object. Do the beetle and robot feel fear, or do they simply express a defensive response? The fact that we have feelings when we act emotional does not mean that every act that looks emotional is accompanied by feelings. The renowned ethologist

Niko Tinbergen reached a similar conclusion: "Although. . . the ethologist does not want to deny the possible existence of subjective phenomena in animals, he claims that it is futile to present them as causes, since they cannot be observed by scientific methods. . . . Hunger, like anger, fear, and so forth, is a phenomenon that can be known only by introspection. When applied to another subject, especially one belonging to another species, it is merely a guess about the possible nature of the animal's subjective states.

Given that we are left with nonverbal behavior as our main tool for assessing emotional states in animals, and that animal studies are the best way to study the brain, how can we escape the credibility problem and gain a richer psychological understanding of the brain's emotional circuits? How can we, in other words, study emotions using nonverbal emotional behavior as a measure without using it as a measure of feelings? Ironically, an important clue about how to do this comes from cognitive science. As we saw in the last chapter, cognitive science was successful because it figured out how to study the mind without getting bogged down in questions about subjective experience. The trick was to treat the mind as an information-processing device rather than as a place where experiences occur. Although early cognitive scientists considered emotions to be more a matter of mental content than of information processing, and thus not subject to cognitive analysis, their processing approach is, in fact, directly applicable to the study of emotion. Just as it is possible to study how the brain processes the color or shape of a stimulus without first figuring out how the conscious experience of color or shape comes about, it is also possible to examine how the brain processes the emotional significance of a stimulus without necessarily first figuring out how that stimulus comes to elicit conscious feelings. Since emotions as processes can be studied in animals and humans alike, and since, as we'll see, emotional processing underlies both emotional behavior and emotional feelings, a processing approach is a way out of the credibility problem.

Even if this method allows us to account for important aspects of emotion in animals without having to resort to subjective states, it should not be taken to mean that subjective states exist only in humans. In the last chapter, I discussed how other animals might have domain-specific forms of consciousness, and in the case of nonhuman primates, domain-independent forms of nonverbal consciousness, but how only humans have verbal working memory, and thus language-based consciousness and the mental frills that language makes possible. The problem is that as soon as we rely on subjective states to explain behavior, we confront our inability to know whether such states really do exist in creatures other than humans. So my approach has been to discuss conscious emotions (feelings) only with respect to humans, and to restrict myself to the notion of emotional processing when I talk about creatures other than humans. In this way, I avoid the construction of a theory that can never be proven, but at the expense of having one that may be incomplete. I'm comfortable with this approach, though others prefer the alternative.

EMOTIONAL PROCESSING

In order to stay alive, remain healthy, and propagate their species, animals have to be able to detect friend and foe, identify safe and nutritious foods, select mates, and respond appropriately to these and other stimuli and situations. Attempting to have sex with rather than defend against a predator would be costly. There must therefore be some mechanism, some circuit, located between input and output systems that translates environmental information into specific responses when certain kinds of stimuli occur. In fact, there must be a number of such circuits in the brain, since different systems are involved in defensive, feeding, and sexual behaviors, to name just a few categories of activities related to survival and well-being. When I use the term *emotional processing*, I have the functions of circuits like these in mind.

Specifically, from this point of view, emotion can be defined as the process by which the brain determines or computes the value of a stimulus. Other aspects of emotion then follow from this computation. First, emotional reactions occur. These overt bodily responses and associated changes in internal body physiology are the advance guard of emotional responsivity. Subsequently (at least in humans), a feeling emerges as we become aware that our brain has determined that something important is present and we are reacting to it. In addition, given that we are in an emotionally arousing situation, we often take action. That is, we do things to cope with or capitalize on the event that is causing us to be emotionally aroused. Emotional actions, in other words, occur when emotions motivate us to do things. I'm going to focus on emotional reactions and feelings here, and save action for the next chapter, which is about motivation.

It's relatively easy to account in neural terms for how emotional reactions follow from emotional processing: information received by sensory systems activates emotional-processing circuits, which evaluate the meaning of the stimulus input and initiate specific emotional responses by triggering output

circuits. Defense, food-seeking, and sex circuits receive inputs from the same sensory systems, and thus receive similar information, but a given circuit is only activated when the sensory influx contains stimulus information relevant to its operation. These detection and reaction processes take place automatically, independent of conscious awareness of the stimulus and feelings about it. This is why credibility is not a problem from the point of view of emotional processing.

The simplest way to illustrate the independence of emotional processing from consciousness in the control of emotional behavior is to describe an example in which consciousness is not a factor. At some point in your life, you've probably jumped out of the way of something rapidly approaching and only afterward noticed what it was—a ball thrown at you or a bus flying by, for example—and also only afterward noticed your heart pumping strongly. The feeling of fear came after you jumped and after your heart was already pumping—the feeling itself did not cause the jumping or the pumping. While an anecdote like this doesn't prove anything, there is also quite a lot of scientific evidence for this kind of reaction. The heroin addict study described above is one example. Let's consider another.

It's possible to present stimuli to the brain subliminally (unconsciously). This can be done in a number of different ways, but one commonly used is backwards masking.²³ In this procedure, an emotionally arousing visual stimulus is flashed on a screen very briefly (for a few milliseconds) and is then followed immediately by some neutral stimulus that stays on the screen for several seconds. The second stimulus blanks out the first, preventing it from entering conscious awareness (by preventing it from entering working memory), but it does not prevent the first from eliciting an emotional reaction (the stimulus still changes the beating of the heart or makes the palms sweat). Since the stimulus never reaches awareness (because it is blocked from working memory), the responses must be based on the unconscious processing of the meaning of the stimulus rather than on the conscious experience of it. By short-circuiting the stages necessary for the stimulus to reach consciousness, the masking procedure reveals processes that go on outside of consciousness in the human brain.

In creatures that lack the kinds of working memory processes that enable information to be held in mind consciously (see the discussion of the evolution of thinking in the last chapter), nonconscious processing is the rule rather than the exception. The basic behavioral repertoire of these creatures is regulated nonconsciously because

they lack the kind of brain that can have a conscious experience. The human brain has these nonconscious capabilities as well. By using behavioral reactions (and associated changes in body physiology; such as responses of internal organs) that occur in emotional situations to index or measure emotional processing in experimental animals, the brain mechanisms that underlie this processing can be uncovered. Studies of people with brain damage, and functional imaging studies of undamaged people, can be used to verify whether the same systems are involved in emotional processing in the human brain (it's much easier to verify that the same system is involved than to make the initial discovery in humans, which is why the animal studies are essential to lead the way). As long as we don't pursue aspects of emotion that are unique to the human brain, and as long as we choose an-

imals that have emotional processes that are relevant to the human emotion we are interested in studying, we can pursue the brain mechanisms of human emotional processing in animal brains.

It's important to distinguish a processing approach to emotion from two other approaches. On the one hand, since a processing approach relies on overt responses, and sidesteps subjective experience, it may sound fundamentally behavioristic. A few years ago, after giving a lecture in which I argued for this approach, I was, in fact, accused of being a radical behaviorist (one who denies that the mind exists) disguised as a neuroscientist. But that description does not suit me for two reasons. First, I'm interested in what goes on in the black box, and especially in how circuits in the brain represent and evaluate the meaning of emotional stimuli. I don't study behavior to understand behavior so much as to understand how processes in the brain work. This interest in internal processes alone disqualifies me from being considered a radical behaviorist. But, also, I'm very interested in feelings and consciousness, topics that behaviorists shunned. While behaviorists turned away from the analysis of the content of subjective states, I, in contrast, want to understand subjective phenomena in terms of underlying processes rather than as conscious content. Later, in fact, I will consider how feelings emerge from neural processes.

A processing approach may also be confused with the so-called cognitive approach to emotions, which treats emotions as appraisals—that is, as thoughts about a given situation. While some appraisal theorists allow for unconscious appraisals (which is consistent with a processing approach), most emphasize appraisals as conscious thoughts and use verbal self-report to understand the nature of the appraisal process. This approach, obviously, takes us right back to the credibility problem. Conscious appraisals may indeed occur during an emotional state, but there are other, more fundamental processes at work as well. An understanding of these more fundamental processes is the goal of the processing approach.

In sidestepping the credibility problem, a processing approach thus allows us to study unconscious emotional functions similarly in humans and animals, and at the same time offers a method of understanding emotional consciousness (feelings) as well (since feelings themselves result from processes that occur unconsciously). In addition, a processing approach presents another advantage. It allows emotion and cognition to be treated the same (as unconscious processes that can but do not necessarily lead to conscious experiences), and it opens the door for the much needed integration of cognition and emotion (and, as we'll discuss later, motivation)—in short, the reassembly of the mental trilogy.

A QUICK FIX?

One way to begin to reassemble the mental trilogy might be to put all the newly acquired information about the thinking brain that came from the cognitive revolution together with the compelling view of the emotional brain provided long ago by the limbic system concept. Perhaps the notion of the limbic system simply needs to be modernized by treating it as an emotional-processing network rather than as the seat of conscious feelings. However, while the limbic system remains the predominant explanation (both in neuroscience and popular culture) of how the brain makes emotions, it is a flawed and inadequate theory of the emotional brain. I made this point very forcefully in *The Emotional Brain*, but those criticisms bear repeating.

The limbic system concept, brainchild of the pioneering neuroscientist Paul MacLean, was put forth in the context of an evolutionary explanation of mind and behavior. It built upon the view, promoted by comparative

anatomists earlier in the twentieth century, that the neocortex is a mammalian specialization—other vertebrates have primordial cortex, but only mammals were believed to have neocortex (chap. 3). And since thinking, reasoning, memory, and problem-solving are especially well developed in mammals, and particularly in humans and other primates that have relatively more neocortical tissue, these cognitive processes were believed to be mediated by the neocortex and not by the old cortex or other brain areas. In contrast, the old cortex and related subcortical regions form the limbic system, which was said to mediate the evolutionarily older aspects of mental life and behavior, our emotions. In this way, cognition came to be thought of as the business of the neocortex, and emotions of the limbic system.

The limbic system theory began to run into trouble almost immediately when it was discovered, in the mid-1950s, that damage to the hippocampus, an old cortical area and the centerpiece of the limbic system, led to severe deficits in a distinctly cognitive function, long-term memory. This finding was incompatible with the original idea that the primitive architecture of the limbic system, and especially of the hippocampus, was poorly suited to participate in cognitive functions. Subsequently, in the late 1960s, it was discovered that the equivalent of mammalian neocortex was present, though in a rudimentary form, in nonmammalian vertebrates (chap. 3). As a result, the old/new cortex distinction broke down, challenging the evolutionary basis of

the assignment of emotion to the old cortex (limbic system) and cognition to the neocortex.

The limbic system itself has been a moving target. Within a few years after its inception, the definition expanded from the original notion of the old cortex and related subcortical forebrain nuclei to include some areas of the midbrain, and even some regions of the neocortex. Several attempts have been made to salvage the limbic system by defining it more precisely. Nevertheless, after half a century of debate and discussion, there are still no generally accepted criteria for stipulating which areas of the brain belong to the limbic system. Some scientists have suggested that the limbic system be abandoned.

In spite of these difficulties, the limbic system continues to survive, both as an anatomical concept and as an explanation of emotions, in textbooks, research articles, and scientific lectures. This is in part attributable to the fact that both the anatomical foundation and the emotional function it was supposed to mediate were defined so vaguely as to be irrefutable. For example, in most discussions of how the limbic system mediates emotion, the meaning of the term *emotion* is not defined. Reading between the lines, it seems that the authors are often referring to something akin to the common English-language use of the term, which is to say feelings. However, as we've seen, a conception of emotion in terms of feelings is problematic.

Further, the anatomical criteria for inclusion of brain areas in the limbic system remain undefined, and evidence that any limbic area, in whatever definition, contributes to any aspect of any emotion has tended to be used as validation for the whole concept. For example, because the amygdala was included as one of the limbic areas, studies showing that the amygdala participates in fear were viewed as evidence that the limbic system theory was correct, in spite of the fact that many other limbic areas played little or no obvious role in fear or other emotions. In spite of many hundreds of experiments aimed at elucidating the role of limbic areas in emotion, there is still very little understanding of how our emotions might be the product of the limbic forebrain. Particularly troubling is the fact that one cannot predict, on the basis of the original limbic theory of emotion or any of its descendants, how specific aspects of emotion work in the brain. The explanations are all *post hoc*, concocted after a given experiment to explain the data—scientists typically put more stock in predictions than explanations. This problem is particularly apparent in recent work using functional imaging to study the human brain. Whenever a so-called emotional task is presented, and a limbic area is activated, the activation is explained by reference to the fact that limbic areas mediate emotion. And when a limbic area is activated in a purely cognitive task, it is often assumed that there must have, in *fact*, been some emotional component to the task. We are, in other words, at a point where the limbic theory has become an off-the-shelf explanation of how the brain works, one grounded in tradition rather than in facts. Deference to the concept is inhibiting creative thought about how mental life is mediated by the brain.

Although the limbic system theory is inadequate as an explanation of the specific brain circuits of emotion,

MacLean's original ideas are insightful and quite interesting in the context of a general evolutionary explanation of emotion and the brain. In particular, the notion that emotions involve relatively primitive circuits that are conserved throughout mammalian evolution seems right on target. Further, the argument that cognitive processes might involve other circuits, and might function relatively independent of emotional circuits, at least in some circumstances, also seems correct. These functional ideas are worth preserving, even if we ultimately abandon the limbic system as an anatomical theory of the emotional brain.

LESS IS MORE

The limbic system theory failed in part because it attempted to account for all emotions simultaneously, and in so doing did not adequately account for any one emotion. When I got involved in emotion research in the late 1970s, I decided to take the opposite approach and study one emotion—fear—in detail. Below, I'm going to focus specifically on what we've learned about fear, because it is the emotion about which we know most. But the basic principles that have been uncovered about the fear system are likely to be applicable to other systems as well. Although different brain circuits may be involved in different emotion functions, the relation of specific emotional-processing circuits to sensory, cognitive, motor, and other systems is likely to be similar across emotion categories.

Much of what we've discovered about the neural basis of fear has come from studies of fear conditioning over the past two decades. This procedure, long a standard tool in behavioral psychology had not been used much to study the brain until I and several other researchers adopted it as way of studying emotional learning circuits. (As mentioned in chapter 5, much of the key work on the neural basis of fear conditioning was done by the labs of Robert and Caroline Blanchard, Bruce Kapp, Michael Davis, and Michael Fanselow, as well as mine; also important were earlier studies by David Co

hen in pigeons.) While some limbic areas did turn out to be involved in fear conditioning, the exact locations and the nature of their involvement would never have been predicted by the limbic system theory alone.

In the end, fear conditioning may not tell us everything we would like to know about fear, especially human fear. For example, the neural circuits involved in responding to conditioned fear stimuli may participate in, but are probably not sufficient to account for, more complex aspects of fear-related behavior, especially responses that depend not on specific stimuli but on abstract concepts and thoughts, such as the fear of failing, fear of being afraid, or of falling in love. Nevertheless, fear conditioning has been an excellent way to start understanding some basic facts about fear, especially how fear responses are coupled to specific stimuli that people and other animals encounter in their daily lives.

What makes fear conditioning so useful from a circuit-tracing point of view? For one thing, it's a simple procedure. All it requires to turn a meaningless stimulus, like a tone pip, into a fear-arousing event is a few occurrences of the pip (often only one) at the same time as an aversive event, like a mild shock to the skin. Also, it's versatile—just about any stimulus that predicts shock or many other kinds of dangerous stimuli can serve as a conditioned fear stimulus. In addition, the learning is long-lasting, and maybe even permanent. And it can be administered similarly in humans and rats, making it possible to study the rat brain for the purpose of understanding human fear. Further, the responses are hardwired and automatic. We don't have to learn to freeze or raise blood pressure in the presence of dangerous stimuli, for the brain is programmed by evolution to do these things. We have to learn what to be afraid of, but not how to act afraid.

Implicit in all this, of course, is a strategy for mapping out the fear-processing circuit. All we have to do is trace the pathway forward from the input system (the sensory system that processes the conditioned stimulus, say the tone pip) to the output system (the system that controls freezing or other hardwired responses). The fear-processing circuits, by this logic, should be located at the intersection of the input and output systems.

NUTS IN YOUR BRAIN

Studies of fear conditioning have shown beyond a doubt that the brain region that sits at the intersection of input and output systems of fear, and the key to understanding how danger is processed by the brain, is the amygdala. Actually, there are two of them, one on each side of the brain, but, for convenience, we'll talk about the amygdala in the singular, since the two do pretty much the same thing.

Once an obscure region of the brain, the amygdala is practically a household word these days. The Batman comic series *Shadow of the Bat* featured a monster called "Amygdala" who was named after the "almond-shaped mass of nerves in the brain that controls feelings of rage" ("almond" is indeed the English translation of *amygdala*, the Greek-derived word used to name this structure because of its appearance). Recently, a newspaper column called "Kids' City" discussed the role of the amygdala in childhood fears. A Website exists where you can "click your amygdala," that is, click on certain buttons to expose yourself to stimuli that will supposedly turn on your amygdala. One night, while channel surfing, I came across a Sci-Fi Channel show in which an alien was able to control people's fears by influencing their amygdalae (that's the plural). I have even been contacted by lawyers hoping to mount an "amygdala defense," an argument that the violent crime performed by their client was the fault of the client's amygdala rather than his free will. For better or worse, the amygdala is no longer obscure. But let's put these issues related to the amygdala's popularity aside for now and go a bit deeper into its synaptic organization and function.

As described in chapter 5, the amygdala contains a dozen or so distinct divisions or areas, of which only two are necessary for fear conditioning. Information about the outside world is transmitted to the lateral nucleus from sensory-processing regions in the thalamus and cortex, allowing the amygdala to monitor the outside world for signs of danger. If the lateral nucleus detects danger, it activates the central nucleus, which initiates the expression of behavioral responses and changes in body physiology that characterize states of fear.

But what makes conditioning occur? We've already considered this in chapters 5 and 6, and will just summarize here. When a sound is presented to a naive animal, it reaches the lateral nucleus and mildly activates neurons there. GABA inhibition prevents much from happening in response, and if the sound is repeated without consequence, the cells quickly stop responding. But if the sound is followed by a shock, the weak preexisting response is greatly amplified following the rules of Hebbian plasticity: the shock activates the postsynaptic cell while the sound is causing the presynaptic terminals to release glutamate. During the activation, calcium enters the postsynaptic cell, and then a host of intracellular chemical reactions, involving kinases and

transcription factors, activate genes, inducing them to make proteins that then stabilize the relation between the presynaptic and postsynaptic neuron. As a result of conditioning, the sound acquires the ability to elicit strong excitation in the amygdala, making it a more potent stimulus for activating amygdala circuits after conditioning than it was before. A stimulus that would normally not get past the GABA guard in the lateral nucleus thereby comes to travel with ease to the central nucleus, where the floodgates of emotional reactivity are opened.

PUTTING FEAR IN ITS PLACE

Clearly, much has been learned about the role of amygdala circuits in fear conditioning. With this information in hand, we can now begin to ask how processing in the amygdala relates to processing in cortical circuits involved in cognitive processing. We start with a consideration of the connections between the hippocampus and the amygdala, and their contribution to the contextualization of fear—that is, the regulation of fear on the basis of our assessment of the situation we are in.

In an emotionally charged situation, there is often some stimulus that stands out, and others that are less prominent, though nevertheless important. For example, if you are being robbed at gunpoint while visiting a foreign city, the most salient factor is the guy with the pistol pointed at you. But the context in which the robbery occurs is also significant—you may well feel uneasy if you go back to the street corner where it happened, or even back to that city or country. (My son and his friends were on their way to a basketball game at another school when one of the boys turned around and went home so that he would not have to cross an intersection where he tripped and broke two teeth several years earlier.) In the laboratory, when a rat is conditioned to a tone-shock combination in a certain

chamber, it will freeze and otherwise act afraid if it simply finds itself back in that chamber at some later point. The tone was the most salient cue, but other cues in the chamber were conditioned as well.

The neural basis of so-called contextual conditioning has been studied extensively in recent years, especially by Russ Phillips in my lab and by Mike Fanselow and colleagues. Like tone conditioning, contextual conditioning is dependent on the amygdala, but unlike tone conditioning, it is also dependent on the hippocampus. Just as the auditory system provides the amygdala with information about a tone, the hippocampus, by virtue of its role in relational/configural/spatial processing (chap. 5), provides the amygdala with information about a context in which the emotional learning is taking place. The context is, in other words, a psychological construction, a kind of memory created on the spot, about the various factors that constitute an emotional situation. And just as tone-shock integration occurs in the amygdala during conditioning, context-shock integration occurs there as well. However, while the lateral nucleus is involved in tone-shock integration, it is not necessary for context-shock integration. The basal nucleus (basal amygdala), which receives the connections from the hippocampus, is critical instead. The basal nucleus communicates with the central nucleus in the control of fear reactions to the context. And by way of these hippocampal-amygdala circuits, fear responses can be adjusted depending on the specifics of the situation: a beast in the wild elicits fear, but one in the zoo just fascinates. Beneath it all, though, is a reactive system that is still prepared to respond, as Charles Darwin found out when he tried, without success, to withhold a response to the strike of a poisonous snake behind a protective glass cage in a zoo.

The idea that the hippocampus is involved in context conditioning has been challenged several times. However, these challenges have so far been met, leaving intact the conclusion that connections between the hippocampus and the amygdala account for contextual processing in fear.

CHECKS AND BALANCES

The amygdala also interacts with the medial prefrontal cortex. This region includes the anterior cingulate and orbital regions, as well as areas in transition between them (infralimbic/prelimbic cortex). These areas send connections to several amygdala regions, including the central nucleus, as well as to brain stem outputs of the central nucleus, allowing cognitive functions organized in prefrontal regions to regulate the amygdala and its fear reactions.

Several years ago, Maria Morgan of my lab pursued the role of the medial prefrontal cortex in fear regulation. She found that the consequences of damage to this region varied, depending on where the lesion was located. Some lesions led to a marked exaggeration of fear reactions—rats with such damage froze much more than controls each time the conditioned fear stimulus (the sound that had been paired with shock) appeared. In contrast, with other lesions, no such exaggeration of fear intensity occurred. However, when the sound was presented without the shock repeatedly until fear reactions no longer occurred, animals with the lesion required many more exposures to the sound than unlesioned animals to extinguish—to stop acting fearful. The role of the medial prefrontal cortex in fear regulation recently has been confirmed in studies by Greg Quirk, my former colleague now in Puerto Rico, and by Rene Garcia in France.

Collectively, this work suggests that the prefrontal cortex and amygdala are reciprocally related. That is, in order for the amygdala to respond to fear reactions, the prefrontal region has to be shut down. By the same logic, when the prefrontal region is active, the amygdala would be inhibited, making it harder to express fear. Pathological fear, then, may occur when the amygdala is unchecked by the prefrontal cortex, and treatment of pathologic fear may require that the patient learn to increase activity in the prefrontal region so that the amygdala is less free to express fear. Clearly, decision-making ability in emotional situations is impaired in humans with damage to the medial and ventral prefrontal cortex, and abnormalities there also may predispose people to develop fear and anxiety disorders. These abnormalities could be due to genetic or epigenetic organization of prefrontal synapses or to experiences that

subtly alter prefrontal synaptic connections. Indeed, the behavior of animals with abnormalities of the medial prefrontal cortex is reminiscent of humans with anxiety disorders: they develop fear reactions that are difficult to regulate. Although objective information about the world may indicate that a situation is not dangerous, because they cannot properly regulate fear circuits, they experience fear and anxiety in these safe situations.

The medial prefrontal cortex may thus serve as an interface between cognitive and emotional systems, allowing cognitive information processing in the prefrontal cortex to regulate emotional processing by the amygdala. In addition, emotional processing by the amygdala may influence decision-making and other cognitive functions of the prefrontal cortex. We'll have more to say about prefrontal-amygdala interactions when we consider the relation between fear circuits and conscious feelings of fear later in this chapter.

OF RATS AND MEN

Bugs, slugs, shellfish, and other invertebrates lack an amygdala. Fear conditioning is taken care of in other ways in these creatures. But for vertebrates (at least reptiles, birds, and all varieties of mammals, including humans), the amygdala is responsible for fear conditioning.

It's much more difficult to study human than rat or other animal brains. However, in the last few years, studies of brain-damaged patients and studies using new brain-imaging techniques have examined the role of the human amygdala in fear and other emotional processes. These investigations still cannot pinpoint the contribution of specific circuits within the amygdala, but have nonetheless been very successful in showing that many of the basic fear-conditioning observations in animals apply equally to the human brain.

The landmark year was 1995, when two studies on the effects of brain damage on human fear conditioning were published. Collaborators of mine (Kevin LaBar, Liz Phelps, Dennis Spencer) working at Yale examined a series of about twenty patients who had undergone unilateral temporal lobectomy, a procedure in which large areas of the temporal lobe (including the amygdala) are removed on one side of the brain in an effort to control severe epilepsy. Regardless of the side of the removal, the patients exhibited impaired fear conditioning. Around the same time, Antonio Damasio and colleagues reported that fear conditioning was disrupted in a patient who had a rare condition that resulted in her sustaining damage restricted to the amygdala on both sides of her brain. In the patients in both our study and theirs, explicit or declarative memory for the conditioning experience was intact, indicating that fear conditioning and declarative memory are separable in the human brain, as in the rat brain. Also, damage to the hippocampus in humans, as it does in rats, disrupts fear conditioning to contextual cues.

These findings on the effects of brain damage on fear in humans have recently been complemented by functional-imaging studies. One study was performed by LaBar, Phelps, and me, and a second was by John Morris, Arne Ohman, and Ray Dolan. Both used functional MRI and both found that the amygdala of humans was activated during fear learning. The Morris study had an additional interesting twist—the visual conditioned stimuli were masked, and thus never reached consciousness, showing that the human amygdala can undergo emotional learning to stimuli that are never experienced.

Some researchers, especially those who work on the neocortex, have argued that the direct pathway from the thalamus to the amygdala, the so-called low road, is unlikely to be very significant in the brains of humans and other primates.^{s0} Their logic is that the cortex is so important in primates that any effect of low-level subcortical processing would be overshadowed. However, a subsequent study by Dolan's group helps put this criticism to rest. Typically, most imaging studies look for areas where neural activity increases or decreases.^{ss} Suppose activity on the average increases in each of two areas, A and B, during three presentations of a stimulus. Suppose further that in area A the activity was high on the first trial, medium on the second, and low on the third. The average change would be a medium increase. The same average increase could occur in area B if the activity was low on the first, medium on the second, and high on the third. This pattern, though, would not support the conclusion that the two areas are functionally connected, since activity did not change the same way in the two areas on each trial. But if the two areas did change in the same way on each trial, then functional connectivity would be implied. Dolan and colleagues did this

kind of analysis on the data they obtained in the masked conditioning study. Specifically, they asked which brain areas changed in a way that would indicate connectivity with the amygdala. They found that across the whole brain amygdala activity during conditioning was most directly related to activity in subcortical visual-processing areas, including an area of the visual thalamus. Particularly significant was the fact that amygdala activity was not at all related to activity in areas of the visual cortex. This finding thus indicates that unconscious emotional learning occurs through the path from visual sensory areas of the thalamus to the amygdala. The low road is indeed used in both the rat and human brain.

Given all the parallels we've seen between the rat and the human, findings from rats should be assumed to apply to humans until they are proven not to, at least in the area of fear reactions. This doesn't mean that the human brain should, like the rat brain, be afraid of cats, but rather that the general wiring plan of the human and rat fear system is the same. As a result, the synaptic circuits that turn on the rat amygdala in the presence of cats will similarly turn on the human amygdala in the presence of stimuli that are dangerous to us.

In addition to studies on the role of the amygdala in conditioning, there have been a number of other imaging studies of humans that are relevant to the topic of fear. For example, it is well known that the expression of emotion on a human face is a potent emotional stimulus. Studies by Dolan's group in London, as well as by Hans Breiter, Paul Whalen, and Scott Rausch, found that exposure of human subjects to fearful or angry faces potently activates the amygdala. Whalen and colleagues found that even masked presentations of such faces cause such activation. Along similar lines, damage to the human amygdala interferes with the ability to judge the emotions expressed in faces and voices. People with amygdala damage, in fact, seem to have trouble in their daily lives deciding whom to trust.

These latter findings in humans are reminiscent of an observation first made in the early 1900s and then popularized in a 1937 report by Kluver and

Bucy: monkeys with temporal lobe lesions that included the amygdala lost their fear of things they were normally quite afraid of, such as humans and snakes. Somewhat later, Edmund Rolls, Taketoshi Ono, and other investigators recorded from neurons in the monkey amygdala in an effort to shed some light on the so-called Kluver-Bucy syndrome. They found cells that responded specifically to faces and other kinds of biologically significant stimuli, like food items and threatening objects.

Recent work in humans by Liz Phelps and Paul Whalen has further implicated the amygdala in social interactions. In separate studies, they found that exposure of white subjects to the faces of unfamiliar African Americans led to amygdala activation, and the degree of activation was directly related to the subjects' score on a test that measures racial biases. Particularly significant is the fact that the bias test was an implicit measure of racial bias. This suggests that implicit (unconscious) tendencies toward racism are reflected in the degree to which the amygdala is activated by stimuli representing the group biased against. This work is taking us into new and provocative areas, but is also raising serious ethical issues for researchers. Given that negative attitudes and biases have their strongest effects on behavior when they are unconscious, and thus cannot be guarded against and compensated for, should researchers inform subjects of these biases? Such studies also force us to confront ethical decisions as a society. How far should we go in using brain imaging to read minds, and how should we use the information we discover? It is testimony to the progress being made that these questions need to be asked.

RELIVING THE EMOTIONAL PAST

Fear conditioning by the amygdala, as I've said many times, is an implicit form of learning, one that does not require conscious participation. During any experience in which we are awake and alert, however, working memory will be aware of what is going on, and if what is going on is significant, the executive will direct the storage of information about the situation in the explicit memory system. We are thus later able to consciously recall (retrieve into working memory) those aspects of the experience that were stored explicitly. While this is true of any kind of explicit memory, explicit memories about emotion are unique.

Explicit memories established during emotional situations are often especially vivid and enduring, and for this

reason are called flashbulb memories. The classic example is that most baby-boomers know where they were and what they were doing when they heard the news that JFK had been shot. But we are all aware from our own daily experiences that we remember particularly well those things that are most important to us, those things that arouse our emotions. Emotions, in short, amplify memories.

Studies by Jim McGaugh and colleagues spanning several decades have implicated the amygdala in the emotional amplification of explicit memory. During emotional arousal, outputs of the central amygdala trigger the release of hormones from the adrenal gland that return to the brain. The amygdala, it turns out, is an important target of such feedback. The feedback consists of the direct action of body hormones (like cortisol released from the adrenal cortex) on amygdala neurons, as well as indirect actions whereby hormones (like epinephrine and norepinephrine released from the adrenal medulla) interact with nerves that travel from the body into the brain and ultimately reach and influence neural activity in the amygdala. By way of its connections with the hippocampus and other regions of the explicit memory system, the amygdala then modulates (strengthens) the consolidation of explicit memories being formed during emotional arousal. Later, the memories are more easily retrieved, and the details of the original experience are more vivid. Thus, in addition to *storing* implicit memories about dangerous situations in its own circuits, the amygdala *modulates* the formation of explicit memories in circuits of the hippocampus and related areas. Studies in recent years led by Larry Cahill and Benno Roozendaal in McGaugh's lab have helped to refine the latter conclusion.

According to the mood congruity hypothesis, memories are more easily retrieved when the emotional state at the time of memory formation matches the state at the time of retrieval." For example, we are more likely to remember sad than happy events when depressed. Perhaps amygdala activation during retrieval facilitates remembrance by re-creating, at least in part, the emotional state (the state of the brain resulting from amygdala activation, and all its consequences, as discussed above) that occurred during the original experience—the more similar the pattern of activation is during learning and retrieval, the more efficient retrieval is likely to be. The unreliability of remembered emotion (Kahneman's remembered utility) and the fallibility of eyewitness testimony may be related to the fact that the emotional state at the time of the retrieval will by necessity be somewhat different from the state at the time of the original experience.

As long as the degree of emotional arousal is moderate during memory formation, memory is strengthened. But if the arousal is strong, especially if it is highly stressful, memory is often impaired. Studies by Robert Sapolsky, Bruce McEwen, Gus Pavlides, David Diamond, Tracy Shors, and Jeansok Kim suggest that stress impairs explicit memory by altering the functioning of the hippocampus. Thus, during highly stressful conditions, the concentration of steroid hormone (cortisol) released from the adrenal cortex rises in the bloodstream. For example, when the stress is induced by threatening stimuli, the amygdala is activated and cortisol is released. The hormone then travels to the brain and binds to receptors in the hippocampus, the net effect of which is to disrupt hippocampal activity, weakening the ability of the temporal lobe memory system to form explicit memories. Stressed rats, as a result, do poorly in tasks that require the hippocampus, such as spatial learning. In addition, it is more difficult to induce hippocampal LTP in stressed rats." If the stress continues, hippocampal cells begin to degenerate and ultimately die. These changes appear to account in part for the memory disturbances typical of stress-related psychiatric conditions such as posttraumatic stress disorder, or PTSD, and depression. Stress hormones also have an adverse impact on the prefrontal cortex, and may contribute to the fact that people often make bad decisions under stress. In contrast to its effects on the hippocampus and prefrontal cortex, intense stress seems to enhance the amygdala's contribution to fear. That stress hormones are involved in this amplification of fear by stress is suggested by studies in which the level of stress hormones in the blood is artificially elevated, which has the same effect as stress itself on fear responses.

In short, the exact conditions that lead to a weakened ability to form explicit memories, and to regulate fear by thinking and reasoning, can also amplify fear reactions and enhance our ability to implicitly store information about stressful or traumatic situations. There's good and bad news here. The good news is that even when the ability to

form explicit memory is impaired, we can store useful information about harmful situations. The bad news is that if we don't know what it is we are learning about, those stimuli might on later occasions trigger fear responses that will be difficult to understand and control, and can lead to pathological rather than adaptive consequences. We'll return to this point when we discuss synaptic sickness in a later chapter.

FEAR ITSELF

Emotional arousal has powerful influences over cognitive processing. Attention, perception, memory, decision-making, and the conscious concomitants of each are all swayed in emotional states. The reason for this is simple: emotional arousal organizes and coordinates brain activity. Here I want to focus on how the emotional coordination of brain activity converts conscious experiences into emotional experiences.

In the last chapter, I argued that our immediate conscious content, the thing we are conscious of at any one moment, is what occupies working memory. If this is correct, then it leads to the conclusion that a feeling (the conscious experience of an emotion) is the representation in working memory of the various elements of an immediate emotional state. In this view, the feeling of being afraid would be a state of consciousness in which working memory integrates the following disparate kinds of information: (1) an immediately present stimulus (say, a snake on the path in front of you); (2) long-term memories about that stimulus (facts you know about snakes and experiences you've had with them); and emotional arousal by the amygdala. The first two are components of any kind of conscious perceptual experience, as the only way to identify an immediately present stimulus is by comparing its physical features (the way it looks or sounds) with memories about the same or similar stimuli. But the third kind of information occurs only during an emotional experience. Amygdala activation, in other words, turns a plain perceptual experience into a fearful one. Using Kahneman's term, the amygdala is the source of instant utility in a threatening Situation.

The key question, then, is: How does the amygdala achieve this alteration of consciousness, this transformation of cognition into emotion, or, better yet, this hostile takeover of consciousness by emotion? The answer, I believe, is that emotion comes to monopolize consciousness, at least in the domain of fear, when the amygdala comes to dominate working memory.

The amygdala can influence working memory in a variety of ways, some of which will be described. The first is by altering sensory processing in cortical areas. Working memory finds out about the outside world from sensory-processing areas, so anything that alters how these areas process sensory stimuli will affect the material that is available to working memory. By way of connections with sensory-processing areas in the cortex, amygdala arousal can modify sensory processing. David Amaral has pointed out that while only the latest stages of sensory processing in the cortex send connections to the amygdala, the amygdala sends connections to all stages, allowing the amygdala to influence even very early processing in the neocortex. That sensory cortex areas are influenced by activity occurring in the amygdala is suggested by Norman Weinberger's studies showing that the rate at which cells in the auditory cortex fire to a tone is increased when that tone is paired with a shock in a fear-conditioning situation together with studies performed by Jorge Armony and Greg Quirk in my lab demonstrating that damage to the amygdala prevents some of the cortical changes from taking place. Because the sensory cortex provides important inputs to working memory, the amygdala can influence working memory by altering processing in the sensory cortex.

The sensory cortex is crucially involved in the activation of the medial temporal lobe memory system (chap. 5). By influencing the sensory cortex, the amygdala can have an impact on the long-term memories that are active and available to working memory. But the amygdala also has strong connections with the rhinal cortex, allowing it to directly influence the medial temporal lobe memory system and thus the memories available to working memory.

The amygdala can also act directly on working memory circuits. Although the amygdala does not have direct connections with the lateral prefrontal cortex, it does have connections with other areas of the prefrontal cortex involved in working memory, including the medial (anterior cingulate) and ventral (orbital) prefrontal cortex. As described in the last chapter, connections within and between these regions constitute distributed circuits that underlie the integrative functions of working memory. Damage to the medial prefrontal cortex in rats, as already discussed,

leads to a loss of fear regulation. And studies of monkeys and humans have implicated the orbital region in processing emotional cues (rewards and punishments) and in the temporary storage of information about such cues. The orbital region is connected with the anterior cingulate, and like the anterior cingulate, it also receives information from the amygdala and hippocampus. Humans with orbital cortex damage become oblivious to social and emotional cues, have poor decisionmaking abilities, and some exhibit sociopathic behavior. In addition to being connected with the amygdala, the anterior cingulate and orbital areas are

intimately connected with one another, as well as with the lateral prefrontal cortex, and each of the prefrontal areas receives information from sensory-processing regions and from areas involved in various aspects of implicit and explicit memory processing. The anterior cingulate and orbital areas thus provide a means through which emotional processing by the amygdala might be related in working memory to immediate sensory information and long-term memories processed in other areas of the cortex.

Attention and working memory are closely related (chap. 7), and recent studies by Liz Phelps have shown that amygdala damage interferes with an important aspect of attention. Normally, if we are attending to one stimulus, we ignore others. This is selective attention, and it allows us to focus our thoughts on the task at hand. But if a second stimulus is emotionally significant, it can override the selection process and slip into working memory. Damage to the amygdala, though, prevents this from occurring. The amygdala, in other words, makes it possible for implicitly processed (unattended) emotional stimuli to make their way into working memory and consciousness.

In addition, the amygdala can influence working memory indirectly by way of projections to the various amine cell groups in the brain stem and forebrain that participate in cortical arousal, including cholinergic, dopaminergic, noradrenergic, and serotonergic systems. In the last chapter, we discussed the importance of dopamine and norepinephrine in the regulation of working memory. And as we saw above, norepinephrine plays a critical role in the amplification of explicit memory during emotional states. These arousal pathways are relatively nonspecific since they influence many cortical areas simultaneously. Specificity comes from the fact that the effects of arousal are most significant on circuits that are active. As a result, if the cortex is focused on some threatening stimulus, the circuits involved will be facilitated by the arousal systems. This will help keep attention focused on the threatening situation.

Finally, once the outputs of the amygdala elicit alarm-related behaviors and accompanying changes in body physiology (fight/flight kinds of responses), the brain begins to receive feedback from the bodily responses. Feedback can be in the form of sensory messages from internal organs (visceral sensations) or from the muscles (proprioceptive sensations) or can be in the form of hormones or peptides released by bodily organs that enter the brain from the bloodstream and influence neural activity. Although the exact manner in which bodily feedback influences working memory is not clear, it is likely that working memory has access to this information in one form or

another. But the feedback from these responses is relatively slow, on the order of seconds, when compared to the feedback that occurs by way of synaptic transmission within the brain, which transpires within a matter of milliseconds. Body feedback adds at least intensity and duration, but may also help refine our interpretation of the emotion we are experiencing, once the episode has been triggered. As we saw above, bodily feedback in the form of stress hormones can either enhance or impair long-term memory functions of the temporal lobe memory system, which will in turn influence the content of working memory.

In the presence of fear-arousing stimuli, activation of the amygdala will lead working memory to receive a greater number of inputs, and inputs of a greater variety, than in the presence of emotionally neutral stimuli. I propose that these extra inputs add affective charge to working memory representations, and are what make a particular subjective experience a fearful emotional experience.

But what about animals without a well-developed prefrontal cortex? Do they have any kinds of emotional experiences? I argued in the last chapter that it might be possible to have certain kinds of modality-specific conscious states when the activity of one system dominates the brain. This might happen with strong sensory stimulation (loud noise or a painful stimulus), or in response to emotionally charged stimuli (sight of a predator). Modality-specific feelings can be thought of in terms of passive states of awareness, as opposed to the more flexible kind of conscious

awareness, complete with on-line decision-making capacities, made possible by working memory.

Although my theory of emotional experience is based on studies of fear, it is meant as a general-purpose theory that is applicable to all kinds of emotional experiences. The particulars will be different, but the overall scheme (whereby working memory integrates sensory information about the immediately present physical stimulus with memories from past experiences with such stimuli and with the current emotional consequences of those stimuli) will apply to all varieties of emotional experience in humans, from fear to anger to joy and dread, and even love.

BEYOND FEAR: THE BRAIN IN LOVE

Whenever I give a lecture on emotions and the brain, the question I am most often asked at the end is: “Is the amygdala involved in emotions besides fear, especially positive emotions?” This is not a question that my work has addressed directly. I’ve been trying to understand specifically the mechanism of fear rather than how the amygdala works in general. But other researchers have studied the role of the amygdala in processing stimuli that predict desirable things (like tasty foods and sexually receptive partners). Included is the work of Barry Everitt and Trevor Robbins at Cambridge University, David Gaffan and Edmund Rolls at Oxford, Norman White at McGill, Michela Gallagher at Johns Hopkins, and Taketoshi Ono at Toyoma University in Japan. However, the implications of this work in animals for understanding specific human emotions are less clear than the relation of findings on fear-processing in animals to human fear. Nevertheless, this research is extremely important and will be a major part of any thorough understanding of how the brain creates emotions. We’ll consider some of this work in the next chapter when we explore how behavior is motivated by emotional processing.

In the rest of this chapter, though, I want to turn to a very interesting set of studies that have begun to approach the neural basis of an emotion that is near and dear to humans—namely, love.

The key to understanding how any mental or behavioral function works in the brain, as I’ve said often in this book, is to be able to study it in experimental animals. Love might therefore seem to be an unlikely topic for brain researchers, one particularly vulnerable to the credibility problem. The key issue is whether there is some way to study the function in animals that makes sense in terms of human behavior. For fear, we were able to use conditioning because conditioned fear responses are similar in humans and other mammals. In the case of love, though, the situation is more complicated for several reasons, not the least of which is the fact that most animals don’t pair up with one another exclusively. Not only did researchers have to find ways to study love behaviorally, they also had to find species that were monogamous.

Only about 3 percent of mammals are monogamous, and even within nonhuman primates, monogamy is fairly rare. One species whose members do remain together as a pair after mating to raise their offspring as a family, even across generations, is the prairie vole, a small rodent living in the midwestern plains in the United States. Given that pair-bonding is so rare, the monogamous prairie vole offers a possible window into the biology of attachment.

Attachment (pair-bond formation) is a key element of love. Perhaps the synaptic mechanisms that underlie attachment in voles are also at work in other species that are fairly monogamous, namely, us.

Vole researchers used a completely different strategy than the one commonly used to study fear. Rather than starting with the circuits and then try-

ing to figure out the chemistry, they started with chemical findings and have attempted to relate them to circuits. Much of the work has been conducted by Tom Insel, Sue Carter, and their colleagues and students. It’s also worth noting that Insel credits his mentor, Paul MacLean, for intellectual inspiration. MacLean’s ideas thus continue to motivate research on the emotional brain more than half a century after his introduction of the limbic system concept.

In a recent article, Insel pointed out two features of prairie voles that made them attractive for studying pair-bonding. The first is that monogamy even occurs in voles living in laboratory settings. If bonding only occurred in the wild, it would be very difficult to study its neural basis. In the laboratory, bonding can be measured by putting a vole in the middle chamber of a box with three compartments. The vole is free to travel to the compartment on either side. In one of these, it encounters its mate, and in the other, a stranger is present. Voles that have mated choose to be with

their partner, whereas unbonded ones have no particular preference. After bonding, if another vole comes into the area, the male engages in mate protection, attacking the intruder.

The second feature of voles that has helped move this work forward is the fact that pair-bonding is present only in prairie voles and not in a closely related creature, montane voles, which are found in the Rockies and live individually rather than in family groups. These animals don't form mate preferences after sexual intercourse, so that when placed in the three-chamber box, they don't spend more time with a vole they mated with than a novel one, and the males don't attack intruders after sex. Differences in the brains of these two kinds of voles might therefore provide important clues about the biology of pair-bonding, family organization, and perhaps love itself

One of the main discoveries was that receptors for two hormones believed to play an important role in reproductive behavior were located in different circuits in prairie and montane voles. These hormones, called vasopressin and oxytocin, are found only in mammals, and are related to ancestral hormones that play a key role in behaviors like nest-building in nonmammalian species. In mammals, they continue to be significant in reproductive behavior. For example, oxytocin is involved in uterine contractions during labor and in milk release during nursing. In the brain, though, these chemicals function not just as hormones, but also as neurotransmitters and/or modulators, being released from nerve terminals and binding to postsynaptic receptors.

For now, we'll skip over the exact location of the circuits in which the chemicals are present, as the relation of the chemicals, circuits, and behavior is still not perfectly clear. What is clear, though, are the functions of these chemicals in the behavioral differences between the voles. This has been determined by injecting drugs into the brain that either stimulate or inhibit the action of vasopressin or oxytocin. The drugs have not been injected into specific brain areas but into the ventricles, cavities that contain CSF (cerebrospinal fluid), which flows from the ventricles into the spaces surrounding neurons. With such injections, the drug will reach widespread areas of the brain and can then influence neural function in those areas that contain cells with appropriate receptors. When a drug that blocks the action of naturally released oxytocin is placed in the ventricles of a female prairie vole just before mating, she mates but does not bond with the sex partner. The drug disrupts attachment, not sex, suggesting that oxytocin released during mating underlies bond formation in females. In contrast, if a drug that blocks vasopressin is placed in the ventricles of a male prairie vole before mating, the male likewise mates but doesn't bond, and also doesn't engage in mate-guarding aggression. But if the same drug is injected after mating, intruder aggression does occur. This pattern of results in males demonstrates that the drug blocks attachment and not sexual or aggressive responses. Thus, blocking oxytocin in female and vasopressin in male prairie voles causes them to behave like montane voles. Oxytocin only affects bonding in female brains, and vasopressin only affects bonding in male brains; the female sex hormone estrogen is key to oxytocin's action, just as testosterone is essential for the normal function of vasopressin.

While oxytocin and vasopressin are also present in the brains of humans and are released during sexual behavior, they have not yet been proven to underlie attachment. Regardless of whether the vole findings on oxytocin and vasopressin end up being completely applicable to the human brain, this work illustrates important principles that will surely guide research for some time.

An important area for future work is to pinpoint the exact circuits in which vasopressin and oxytocin are acting during sex-related pair-bond formation. Many studies have been performed on the neural basis of sexual behavior, especially in rodents. Implicated are areas within the amygdala (medial and posterior nuclei), the so-called extended amygdala (an extension of the amygdala into the bed nucleus of the stria terminalis), the striatum (especially the nucleus accumbens), and the hypothalamus (including ventro

medial, medial preoptic, paraventricular, and supraoptic areas). Given that these regions are intimately connected and that receptors for oxytocin and vasopressin are present on neurons in several of them, it is tempting to think of them as forming synaptic circuits crucial not just for sexual behavior, but also for pair-bonding, but this remains unproven.

Areas of the amygdala are included in both the fear and sex circuits. However, the circuits are otherwise quite distinct. Even within the amygdala, different areas are involved in sex (medial and posterior nuclei) and fear

(lateral and central nuclei). This finding emphasizes the importance of mapping the circuit for different kinds of emotional systems rather than assuming that there is a universal circuitry for all emotions. At the same time, different emotion circuits, like the fear and sex circuits, sometimes interact with one another. For example, the medial nucleus sends connections to the central nucleus, where oxytocin receptors are present. This may be related to the ability of both oxytocin and positive social interactions to reduce fear and stress.

Pair-bonding in animals has provided researchers a way of studying something akin to love without having to confront the credibility problem that inevitably arises when any emotion is considered. But in the end we want to know not just about attachment behavior but also about the particular feelings of love. Although we have little research to draw upon at this point, we can use our more detailed understanding of cognitive-emotional interactions in fear to speculate about how the brain does feel love.

Suppose you unexpectedly see a person you care about. Suddenly, you feel the love you have for that person. Let's follow the flow of information from the visual system through the brain to the point of the experience of love as best we can. First of all, the stimulus will flow from the visual system to the prefrontal cortex (putting an image of the loved one in working memory). The stimulus also reaches the explicit memory system of the temporal lobe and activates memories about that person. Working memory then retrieves relevant memories and integrates them with the image of the person. Simultaneously with these processes, the subcortical areas presumed to be involved in attachment will be activated (the exact paths by which the stimulus reaches these areas is not known, however). Activation of attachment circuits then impacts on working memory in several ways. One involves direct connections from the attachment areas to the prefrontal cortex (as with fear, it is the medial prefrontal region that is connected with subcortical attachment areas). Activation of attachment circuits also leads to activation of brain stem arousal networks, which then participate in the focusing of attention on the loved one by working memory. Bodily responses will also be initiated as outputs of attachment circuits, and contrast with the alarm responses initiated by fear and stress circuits. We approach rather than try to escape from or avoid the person, and these behavioral differences are accompanied by different physiological conditions within the body. This pattern of inputs to working memory from within the brain and from the body biases us more toward an open and accepting mode of processing than toward tension and vigilance. The net result in working memory is the feeling of love.

This hypothesis is probably incomplete, but it is probably not completely wrong. It shows how we can build upon research on one emotion to generate hypotheses about others. Given that so much of who we are is defined by our emotions, it is important that we uncover as much as we can about the brain mechanisms of many emotions. This task is just beginning, but the future is bright.

CHAPTER NINE

THE LOST WORLD

EVERY WHY HATH A WHEREFORE.

—Shakespeare, *The Comedy of Errors*, act 2, scene 2

Why do we do the things we do, think the things we think, and make the decisions we make? *Why* questions—questions about our motivations—are fundamental to understanding what makes each of us unique. Following the cognitive revolution, though, motivations, like emotions, were overlooked by many psychologists and brain scientists. The rehabilitation of motivation in both these fields has been slower than that of emotion, but it is also under way. And, as with emotion, the pressure to revive research on and theory about motivation is coming as much if not more from brain science as psychology. So let's rediscover the lost world of motivation, part three of the mental trilogy, from the perspective of how it works in the brain.

ACTING OUT

On a warm summer evening in 1996, a crowd was enjoying a concert in Olympic Park in Atlanta. Suddenly, a bomb exploded. The scene, captured on video, was shown repeatedly on CNN. As soon as the explosion occurred, nearly everyone in the crowd reacted by freezing. After remaining motionless for several seconds, some people starting running away, and soon everyone was in motion.

This short segment of video is a good demonstration of the way an emotional episode unfolds over time. In situations of sudden danger, we initially react by using evolutionarily programmed responses, like freezing, that have long been successful in keeping organisms like us alive. These reactions are elicited, not emitted willfully, and occur automatically, before we have time to think. But we can't freeze forever; sooner or later, we have to take action. In the last chapter, we considered how emotional reactions are automatically triggered by external stimuli. In this chapter, we focus on how emotions, once activated, motivate us to do things, to take actions.

Motivation has many definitions. I use the term here to refer to neural activity that guides us toward goals, outcomes that we desire and for which we will exert effort, or ones that we dread and will exert effort to prevent, escape from, or avoid. Goals direct action, and can be as concrete as a specific stimulus (for example, a particular consumer product) or as abstract as a belief or idea (for example, the belief that hard work will lead to success, or the idea that freedom is worth dying for).

Goal objects are also called incentives. Some are intrinsically motivating (as in the case of food, water, and painful stimulation), while others acquire motivating properties through our experiences with them. In the latter case, they are called secondary incentives. These can arise by association (when a stimulus with low value occurs in connection with one of higher value, as in classical conditioning), by observational learning (seeing the way a stimulus affects other people), by word of mouth (hearing about whether something is good or bad), or by sheer force of imagination.

The view of motivation that I'm going to pursue here is that incentives do their motivating by activating emotion systems. Freezing in reaction to a bomb reflects the activation of an emotional system, whereas running away after several seconds of freezing reflects the motivational consequences of that activation. Actions motivated by emotional arousal have a purpose—to deal with the emotion being aroused. Whether all motivated actions are

necessarily based on emotional activity is debatable. But that emotions are powerful motivators seems indisputable.

Goal-directed behavior is best thought of in functional terms. While we normally try to escape from or avoid harmful stimuli, in some cases, we have to actively engage a dangerous object in order to achieve protection. With motivated behavior, especially in humans, it often is the relation of the act to the goal, not the act itself, that is key. In fact, for humans, most motivated acts are only arbitrarily related to particular goals. When hungry, we can obtain food in many different ways—walking or driving to a restaurant or a store, calling for a delivery, asking a friend to pick up something. Each type of activity is appropriate under certain circumstances, but none has a predetermined, inevitable relation to food acquisition and consumption. The same general kinds of activities (walking, driving, calling, asking) can be used to achieve many varied kinds of goals. Figuring out what to do under different sets of circumstances in order to achieve your goals is what life is all about. Learning and motivation are thus closely intertwined topics.

PUSH AND PULL

My interest in the topic of motivation was sparked by seeing the freezing crowd take flight in the video of the Olympic bombing. This behavior reminded me of what happened in a classic experiment performed by Neal Miller in the 1940s.'

Miller placed rats in an apparatus that consisted of two compartments, one painted white and the other black, separated by an open doorway. The rats started out in the white area, where they received a shock after a delay. At first, they simply froze. But then they figured out that they could escape from the shock by going into the black area. Eventually, they learned that they could avoid the shock altogether by running out of the white compartment as soon as they were placed in it. The rat came to routinely—in fact, habitually—perform the response, even if the shock was turned off. Then, in a later phase of the study, the door between the two areas was closed, rendering the old habit no longer useful. But through trial-and-error behavior, the rats learned that if they turned a small wheel, the door opened. They thus learned a new habit.

Miller interpreted his results in terms of the predominant theory of motivation at the time, Clark Hull's drive theory. In Hull's view, all learning involves the reduction of basic drives (like hunger, thirst, sex, or pain), and current behavior is therefore a product of drive reduction in the past. That is, what we do today in a certain situation is a function of what we did in the past that was successful in reducing drives in similar situations. This is essentially a psychological version of the philosophical position known as hedonism, the idea that we live our lives in such a way as to seek pleasure and avoid pain.

In the late nineteenth century, the pioneer experimental psychologist Edward Thorndike used hedonism as the basis for a psychological methodology. Thorndike showed that hungry cats would learn complex behavioral responses that had no natural relation to feeding but that allowed them to open a door and obtain a bite of food that was visible outside the cage. The cats basically went through a process of trial and error, and those responses that opened the door were then later repeated. Thorndike called this the law of effect: behaviors that are effective in obtaining desirable goals and avoiding undesirable ones are rewarded or reinforced and then repeated, while those that fail to obtain desirable goals or that lead to undesirable ones are punished and are not repeated.

The kind of learning task used by Thorndike came to be called instrumental conditioning—referring to the fact that the behavior is instrumental in achieving the reward or punishment. Because the rewarding or punishing stimulus is associated with the behavioral response, instrumental condition-

ing is called stimulus-response learning. Instrumental conditioning contrasts with Pavlovian or classical conditioning, where the rewarding or punishing stimulus occurs regardless of what the animal does. The association formed is thus not between the rewarding or punishing stimulus and a response, but between neutral stimuli and the reward or punishment they occur with and predict. Hence, Pavlovian conditioning is called stimulus-stimulus learning.

For example, Pavlov's dog salivated to the sound of the bell because that sound had been previously associated with food. The dog did nothing to obtain the food, which simply appeared when the bell sounded (this is a stimulus-stimulus association). On the other hand, if the dog had learned that walking over to a certain spot and

pressing a lever with its paw when the bell sounded would lead to a real piece of food, that behavior would be an instrumental response, rewarded by its success in obtaining food (this would be a stimulus-response association). Together, these two forms of learning became the foundation of behavioral psychology, and behaviorists like John Watson, B. F. Skinner, and Hull hoped to explain all of human behavior in such terms.

Hull's particular contribution to the field was to reinterpret the law of effect in terms of drive reduction, arguing that new instrumental behaviors, new habits, are learned and repeated because they reduce drives. But for Hull, drive only activated or aroused behavior; it did not direct behavior toward specific stimuli that would reduce the drive. Behavioral direction in the presence of drive-elicited behavioral activation was instead based on learned habits, responses that had been successful in reducing drives in previous circumstances.

In this framework, what was the drive that was reduced in the last part of Miller's study, where there was no shock? What moved the rats to learn the new response? Being a drive theorist, Miller assumed some drive must have been reduced, since learning occurred. Because the shock was never presented during the last phase of learning, however, pain prevention could not have been the drive being satisfied. Miller argued that the drive being reduced must have been fear. Fear, in his view, was a learned or acquired drive, as opposed to a biological one (that is, an unlearned, innate or instinctual drive).

The notion of acquired drive added tremendous flexibility to drive theory, since most human behavior is not motivated by pain or deprivation of essential nutrients. Money, for example, has no intrinsic value, and acquires its value only from a mutual acceptance of its worth by those who use it. Similarly, words of praise or scorn are not innately motivating, and achieve their results only by convention.

But drive theory; even with Miller's clever addition, remained problematic for a variety of reasons: One of the main difficulties came from studies showing that rats would learn maze problems to obtain saccharin. Saccharin has incentive value (rats will work to obtain it) due to its sweet taste, but it is non-nutritive and thus cannot satisfy hunger and reduce the supposed drive activated by the biological need for food. As a result of this and other research, drive theory eventually was replaced by the incentive theory. While drives push us from within, incentives pull us from without.

The advantage of an incentive view is that it avoids the postulation of a hypothetical drive state that has to be reduced in order for learning to occur. While incentive theories are therefore less cumbersome, they have their own shortcoming—they often replace the drive problem with the credibility problem (chap. 8). That is, they assume that subjective hedonic experiences, emotional feelings, are what motivate behavior.

The idea I want to develop here is that motivation can be thought of in terms of incentives without assuming that feelings are necessary to translate incentives into actions. I believe that all we need to accept is that in the presence of conditioned (learned) or unconditioned (innate) incentives, emotion systems are activated, placing the brain in a state where an instrumental response becomes a highly probable outcome. In this view, we don't have to postulate the existence of hypothetical concepts like drives or subjective states to explain motivated action. All we need to talk about is real brain systems and their functions.

The brain can be thought of as having a variety of systems that it uses to interact with the environment and keep itself alive. I've used the term *emotion systems* to characterize many of these, but what's relevant here is not the label we give them but the particular function they perform. Included are systems that detect and respond to predators and other dangers, to sexual partners, to suitable food and drink, to safe shelter, and so on. The stimuli (incentives) that activate these systems do so either because of biologically predetermined factors or because of past learning by the individual. When the brain is activated by either an unconditioned or a conditioned incentive, animals, including people, are motivated to perform instrumental responses. These emotionally primed instrumental responses have as their goal, their motive, the alteration of the brain state, the emotional state, that the organism is in.

So long as we resist the temptation to think of emotional states in subjective terms, and instead conceive of them in terms of brain states, we remain on solid ground and avoid the perilous credibility problem discussed in the last chapter.

FEARFUL HABITS IN THE BRAIN

In the years following Miller's classic studies, researchers used aversive instrumental conditioning procedures to study the brain mechanisms of fear.

In such studies, rats or other animals learn to do things that avoid the delivery of a shock. Once they learn how to avoid the shock, they perform the response habitually. Given that the hallmark of fear or anxiety disorders is habitual avoidance of situations that might lead to harm or anxiety, this approach seemed to offer the opportunity to understand the neural basis of a clinically relevant kind of learning. However, as mentioned previously, studies of avoidance conditioning failed to lead to a clear understanding of the neural basis of fear. Little effort was made to distinguish the contribution of neural systems to the two kinds of learning that take place in these tasks—initially, the subject undergoes classical fear conditioning, where cues in the apparatus come to be associated with the shock, then an instrumental avoidance habit is learned on the basis of its ability to remove the animal from the situation in which shock is likely." The notion of multiple memory systems had not yet emerged, and it wasn't fully appreciated that different kinds of learning involve different brain systems. Once researchers started focusing on fear conditioning on its own, stripped out of the context of avoidance conditioning, progress was swift. And as the popularity of fear conditioning increased, the more complex avoidance-conditioning procedures, which were more difficult to relate to brain mechanisms, fell out of favor.

But fear conditioning tells us only about how fear reactions work. If we want to understand how the arousal of an emotion system motivates behavior, we need to turn to instrumental tasks, like avoidance, where the brain learns to take action. With that in mind, after seeing the Olympic bombing video, and being reminded of Miller's rats first freezing and then escaping, I decided several years ago to revisit the question about how fear-motivated instrumental learning—fear habit learning—works in the brain. I didn't, however, want to study just any fear habit-learning task, but one that would enable us to build upon our extensive knowledge of the neural basis of conditioned fear, especially conditioned fear elicited by a tone previously paired with a shock.

Karim Nader, Prin Amorapanth, and I chose to study this using a task called escape from fear. In this procedure, the rats first underwent fear conditioning to a tone-shock combination. As in our other studies, we used the extent to which they froze in the presence of the tone as an index of learning in the fear-conditioning phase. The next day, the animals were placed in a new chamber. After a delay, the tone came on. At first, they froze to the tone. But, eventually, they moved around a bit, and through trial and error, they discovered that movement to the other side of the chamber terminated the tone. Over many trials, the rats learned to escape to the other side immediately to minimize the exposure to the sound. The rats thus learned to take action on the basis of having first learned that tone was associated with shock. The shock, by the way, was never presented in the new chamber. The tone had become a negative secondary incentive, and behaviors that eliminated this stimulus were reinforced and learned as a habit. The design of this experiment was very similar to that of Miller's study, except that we used a tone as the fear-arousing stimulus rather than contextual cues in the apparatus itself. We were now ready to explore the brain mechanisms through which fear system activation motivates action.

Recall from earlier chapters that in order for tone conditioning to occur, the tone has to be transmitted to the lateral nucleus of the amygdala. This region then sends signals to the central nucleus, which controls fear responses. So the first question we asked was whether damage to the lateral or central nucleus would also disrupt the conditioning of the instrumental escape response. The results were clear. Damage to the lateral nucleus prevented the learning of the escape response, but central nucleus lesions had no effect. Thus, the outputs of the central nucleus to the brain stem, required for fear reactions, are not used to learn to do new behaviors, to take action, on the basis of fear system activation. We therefore asked whether damage to one of the other targets of the lateral nucleus—namely, the basal nucleus—would have an effect. This lesion interfered with the ability of the rats to use the information learned during fear conditioning to initiate protective action. It's important to note that basal nucleus damage

had no effect on the animal's ability to freeze to the tone in the first part of the study. This pattern of suits is what neuroscientists call a double dissociation, where the roles of two brain areas in two behaviors are shown to be distinct: the central nucleus was found to be involved in fear reaction but not fear action, and the basal nucleus in the fear action but not reaction.

So exactly how does a conditioned stimulus, a conditioned incentive, processed by the amygdala motivate and reinforce behavior? To answer this question, we need to consider the nature of reinforcement further.

JUST REWARDS

In the early 1950s, studies by Jim Olds and Peter Mimer, two researchers in Donald Hebb's department at McGill, found that rats would return to the location in a chamber where they received a burst of electrical pulses to their brain. The animals, in other words, seemed to be coming back for more. Olds and Mimer knew right away that they were on to something. The stimulation stamped in the behavior, just as rewards do. At the time, behaviorism still dominated psychology, and brain stimulation reward seemed to be a way to figure out the brain mechanisms underlying the topic most important to behaviorism: how responses that lead to rewards are learned.

Olds and Mimer devised a way to test whether brain stimulation could reinforce new responses. Electrodes were implanted in the rats' brains and connected to a lever that the subjects themselves could press. Each time they pressed, they got a few pulses. The rats began to press like crazy. This so-called self-stimulation was clearly rewarding—it motivated the acquisition of an arbitrary behavior (pressing a bar).

This fascinating result was due to a mistake. The researchers were trying to stimulate the reticular formation in the brain stem, a region involved in arousal, alertness, and vigilance. This was in the heyday of drive theory, which proposed that drives lead to learning by arousing the organism, and Olds and Milner wanted to investigate whether an increase in arousal, produced by reticular formation stimulation, would enhance learning. The electrode, though, accidentally ended up somewhere in the forebrain rather than in the brain stem. Although the exact area that the electrode was in is not known, in subsequent studies, Olds and Milner and others identified many sites from which self-stimulation could be elicited, the most potent of which were in the hypothalamus. The reason these sites were so effective was not because the hypothalamus is the reward center of the brain, but because a major nerve pathway passes through the hypothalamus. This pathway, called the medial forebrain bundle, was actually the source of the rewarding effect, as I'll explain later.

It wasn't long after the phenomenon of brain stimulation reward was discovered that the idea arose that its effects were due to the stimulation of "pleasure centers" in the brain. This notion was further fueled when a surgeon in New Orleans, Robert Heath, reported that schizophrenic patients found such stimulations pleasurable. Around the same time, Michael Crichton's literary career was jump-started when he popularized pleasure centers in *The Terminal Man*. While many researchers also treated brain stimulation reward in terms of subjectively experienced pleasure, Peter Shizgall, a leading thinker in this field, has argued that the ability of rewards to motivate behavior and to give rise to pleasurable feelings are separate. This is a motivational version of my notion that emotional behaviors are not necessarily caused by emotional feelings, as Shizgall pointed out.

In spite of its immense popularity, brain stimulation reward research eventually ran out of steam. One reason was that the motivational nature of brain stimulation reward was never quite figured out. Did it activate drives, enhance incentives, or both? Was it the same as natural reward, and could it explain learning? These issues remained unresolved, and by the late 1960s, brain stimulation reward, like other topics related to motivation and emotion, was dying out as the influence of cognitive science was increasing. Drives, incentives, and rewards were just not as important to cognitive scientists as they had been to behaviorists.

PLEASE RELEASE ME

Perhaps the most significant outcome of brain stimulation research was the discovery of why medial

forebrain bundle stimulation was rewarding.” When this pathway is stimulated in the region of the hypothalamus, fibers headed from the forebrain to the brain stem are activated. The main targets of these fibers are neurons that make dopamine,³⁰ which are located in a region of the brain stem called the ventral tegmental area. These cells, in turn, send their axons throughout the forebrain. As a result, when the dopamine cells are activated by inputs from the medial forebrain bundle, they release dopamine widely in the forebrain.”

Dopamine has long been believed to be a critical factor in reward processes. Although there are rewarding conditions that do not depend on dopamine,” much of what we know about rewards centers around the role of dopamine. For example, treating rats with drugs that block the effects of dopamine at its receptors in the brain eliminates the rewarding effects of brain stimulation— that is, rats are much less inclined to press to get the brain jolts under such treatment. Further, if hungry rats are given food in one compartment of a two-chambered apparatus, or if satiated rats are given rewarding brain stimulation in one of the compartments, they will later spend more time in that compartment. This is called a place preference. Treatment of rats with drugs that block the action of dopamine prevents the formation of the place preference. A place preference can also be established by giving rats a shot of amphetamine or cocaine, both of which mimic the action of dopamine at its receptors. It’s no accident that these widely abused drugs work like rewards in learning situations, and the relation of dopamine to drug addiction was important in sustaining interest in reward and motivation during the years of cognitive domination in neuroscience.

Just as brain stimulation reward was initially thought to be due to activation of pleasure centers, dopamine was believed to be the chemical of pleasure. However, as we’ve seen, the hedonistic (subjective pleasure) view of brain stimulation reward is incorrect, and the hedonistic interpretation of dopamine’s role in reward is incorrect as well. For example, blockade of dopamine interferes with instrumental responses motivated by a sweet reward but does not alter the actual consumption of the tasty stuff when it is obtained—the animals still “like” the reward when they consume it, but they are no longer motivated to work for it. Dopamine is thus more involved in anticipatory behaviors (looking for food or drink or a sexual partner) than in consummatory responses (eating, drinking, having sex). But being hungry or thirsty is unpleasant. Pleasure, to the extent it is experienced (see discussion of the credibility problem in chapter 8), would not come during the anticipatory state but instead during consummation. Since dopamine is involved only in the anticipatory phase, and not in the consummatory phase, its effects (at least in the case of primary need states) cannot be explained in terms of pleasure.

The exact role of dopamine in motivation, reward, and habit learning is still being debated. Although it seems clear that dopamine is not involved in subjective pleasure or in the expression of consummatory responses, there is less agreement about the conditions that depend on dopamine. Some adhere to the classic hypothesis—that it is the basis of reward. Another view is that dopamine release is important for the initiation and maintenance of anticipatory behaviors in the presence of secondary incentive stimuli. Others argue that dopamine release notifies the forebrain that something novel or unexpected has occurred, but not that reward per se has occurred. Still others propose that dopamine is involved in the switching of attention and selection of action. These are not mutually exclusive views, and in fact each appears to correctly characterize certain aspects of what dopamine contributes to motivation.

A MOTIVE CIRCUIT

While activation of dopamine neurons in the ventral tegmental area leads to the release of dopamine in many parts of the forebrain, an area called the nucleus accumbens is particularly germane to reward and motivation. Many of the effects of dopamine-related drugs described above can be achieved by applying the drugs directly to the nucleus accumbens, a region of the stria

turn located in front of the amygdala near the bottom of the forebrain. For example, animals will press a bar to administer dopamine or related drugs (cocaine or amphetamine) into the nucleus accumbens. Also, dopamine levels rise in the nucleus accumbens in response to natural rewards (food, water, and sexual stimuli), conditioned incentives (stimuli associated with rewards), and brain stimulation reward. Finally, blockade of dopamine receptors in the nucleus accumbens greatly reduces the rewarding effects of medial forebrain bundle stimulation and of natural

rewards, and also prevents the development of place preferences.

So how, precisely, does an elevation of dopamine in the nucleus accumbens accomplish all of this? More than twenty years ago, Ann Graybiel and Gordon Mogenson suggested that the nucleus accumbens sits at the crossroads of emotion and movement, and that dopamine release in this region plays a crucial role in motivated or goal-directed behavior. This conclusion was based on four main sets of observation. First, the nucleus accumbens receives massive dopamine inputs from the tegmentum. Second, injection of amphetamine or cocaine into the nucleus accumbens leads to behavioral activation—animals start exploring their environment as if in search of something. Third, the accumbens also receives inputs from areas involved in emotional processing, such as the amygdala. Fourth, the accumbens sends output to areas involved in the control of movement (such as the pallidum, an area that connects with the movement-control regions in the cortex and brain stem). Today, it is widely accepted that the nucleus accumbens and areas with which it is connected constitute key elements of a circuit through which emotional stimuli direct behavior toward goals. Let's consider the function of this motive circuit in the broader context of emotional information processing by the brain.

In the presence of an emotionally arousing stimulus, the brain is placed in a state, sometimes called a motive state, that leads to coordinated information processing within and across regions, and results in the invigoration and guidance of behavior toward positive goals and away from aversive. Most of what we know about how incentives are learned and reacted to comes from studies of aversive conditioning, but most of what we know about how conditioned incentives are used to invigorate and guide behavior comes from studies of positive motivation. I'm therefore going to apply to negative motivation what is known about the role of the accumbens in positive motivation in an effort to move forward from our understanding of aversive conditioning into motivation circuits (the details may therefore need to be revised after more research is performed on the role of the accumbens in negative motivation).

The processing begins in the sensory system that receives the stimulus (the auditory system, in the case of a tone that's been paired with an aversive shock) and then flows forward to the amygdala, especially the lateral nucleus, which, in turn, activates the central nucleus of the amygdala (chaps. 5 and 8). Outputs of the central nucleus initiate the expression of species-typical defense responses (like freezing and associated autonomic changes) as well as activate arousal systems in the brain stem, including the dopamine neurons in the ventral tegmental area. The tegmental cells then release dopamine from their axon terminals in the nucleus accumbens (as well as in other areas of the forebrain).

As noted above, animals become active or invigorated when dopamine is injected into the accumbens. This occurs because dopamine facilitates synaptic transmission in the pathway from the accumbens to the pallidum, which in turn connects with movement-control regions in the cortex and brain stem. With the pallidal output amplified, the motor regions are strongly activated, and movement is initiated. Behavior can potentially be invigorated by anything that activates tegmental cells and causes them to release dopamine in the accumbens. Novel stimuli, and conditioned and unconditioned incentives, are prime examples of invigorating stimuli.

But invigoration alone is not sufficient: behavior also needs to be guided or directed. Guidance of behavior is the job of conditioned incentives processed by the amygdala. The basal nucleus of the amygdala, as we've seen, receives information about a conditioned incentive from the lateral nucleus. It then transfers this conditioned incentive to the accumbens. When dopamine is elevated in the accumbens (as a result of the central nucleus's activation of dopamine neurons in the tegmentum that release dopamine in the accumbens), the arrival of an incentive stimulus in the accumbens from the basal nucleus will have a bigger effect on the activity of accumbens cells, and presumably on neurons in the ventral pallidum that are downstream from the accumbens. The incentive thus leads to the release of dopamine, and dopamine facilitates the ability of the incentive to both invigorate and direct behavior.

Conditioned incentives are secondary reinforcers and can stamp in responses that lead to them. Animals will work to obtain conditioned incentives, even when they do not lead to primary incentives. As we've seen, after learning that a tone predicts shock, animals will take actions to stop the tone. Similarly, animals will perform a task (press a bar) to turn on a tone that was previously predictive of some tasty food, even if the tone no longer leads to food (this is

comparable to looking at pictures of a loved one when he or she is away).

The manner in which conditioned incentives reinforce responses is not fully understood. However, LTP occurs in accumbens circuits, and dopamine is essential for the synaptic changes to occur. A reasonable hypothesis is that dopamine facilitates Hebbian plasticity between active presynaptic and postsynaptic cells and thereby strengthens transmission between accumbens pathways that process the incentive and that control responses.

My conception of the role of the amygdala and accumbens in motivation borrows heavily from the work of Barry Everitt and Trevor Robbins at Cambridge University. These researchers have extensively studied how interactions between the amygdala and nucleus accumbens contribute to motivation, especially motivation by positive incentives. They've shown that damage to the basal amygdala eliminates the ability of conditioned incentives to facilitate the learning of a new response by serving as a secondary reward. This disruption occurs for conditioned incentives created by pairing neutral stimuli (like tones) with a variety of primary rewards, including tasty food items,

drugs like cocaine and amphetamine, and exposure of a male rat to a sexually receptive female. Through a clever combination of lesions, where the basal amygdala is lesioned on one side of the brain and the nucleus accumbens on the other, Everitt and Robbins have shown that connections between the amygdala and accumbens are key to the ability of a conditioned incentive to motivate new learning. Their work is part of a long tradition of research on the role of the amygdala in stimulus-reward learning. Some of the other key researchers in this area are Mort Mishkin, Edmund Rolls, Taketoshi Ono, Norman White, Michela Gallagher, and David Gaffan."

Once an emotional habit is well learned, the brain systems involved in expressing it become simpler. The amygdala, for example, drops out of the circuit. After you know how to successfully avoid a specific danger, you no longer need the amygdala, because fear is no longer aroused. A dog needs its amygdala to learn that playing around in the road is dangerous, but once the learning has occurred, he can happily play in the yard next to the road. (In this case, avoidance of danger doesn't arouse fear, it prevents fear.) Characteristically, amygdala-dependent signs of emotional arousal, such as elevations of heart rate, occur during the initial phase of avoidance learning but disappear as the avoidance response is learned. The accumbens likewise drops out once the response is learned; while it's needed to do the learning, it is not necessary to perform well-learned responses. The exact systems that take over in learned defensive habits are not fully understood. Perhaps the accumbens, like other areas of the striatum, trains prefrontal cortical circuits, especially areas of the motor cortex, how to respond. The learning, in other words, might be transferred from the striatum (accumbens) to the cortex. This would be similar to what occurs in declarative or explicit learning: initially, both the hippocampus and neocortex are involved, but once the hippocampus has slowly taught the neocortex the memory, the memory persists without the aid of the hippocampus.

Interactions between the accumbens and amygdala thus go a long way toward accounting, in neural terms, for some of the key aspects of motivation. But the motive circuit consists of other areas as well. The hippocampus participates by way of its connections to the amygdala and accumbens, and may well be involved in the guidance of behavior on the basis of spatial and other kinds of relational cues in the environment: in order to find good things and avoid bad ones, you have to know where you are, where you need to go, how to navigate from where you are to where you need to go, and what sorts of stimuli along the way will be useful in guiding you. This information, drawn from long-term memory, is essential. It is also important to keep the overall goal in mind, and to constantly update your progress toward the goal. Working-memory functions of the prefrontal cortex probably are important for this ability. The prefrontal cortex receives dopamine inputs, and is connected with the accumbens, amygdala, and hippocampus. Certainly, when motivation is based on decisions, the prefrontal cortex, as we will now see, will often be involved.

DECISIONS UNDER FIRE

The study that Nader, Amoranpanth, and I conducted, in which rats learned to escape from a sound that had been associated with a shock, was designed to address the question of how emotional arousal motivates action. The approach we took involved habit learning. But in real-life situations, we don't always have habits to fall back on, or

time to learn a new one. We aren't only impelled by drives, pulled by incentives, and shaped by reinforcers. We live in a complex world where the physical and social environment changes from moment to moment and we often integrate immediate needs and past learning with predictions about the best course of action to take. We use our capacity to think, reason, and evaluate. We make decisions.

Decision-making compresses trial-and-error learning experiences into an instantaneous mental evaluation about what the consequence of a particular action will be for a given situation. It requires the on-line integration of information from diverse sources: perceptual information about the stimulus and situation, relevant facts and experiences stored in memory, feedback from emotional systems and the physiological consequences of emotional arousal, expectations about the consequences of different courses of action, and the like. This sort of integrative processing, as we've seen, is the business of working memory circuits in the prefrontal cortex. In chapters 7 and 8, we discussed the role of the prefrontal cortex in working memory and considered the contribution of the lateral and medial prefrontal cortex. Here, we will focus on two of the subareas of the medial prefrontal cortex in light of their relation to the motive circuits outlined above.

The anterior cingulate cortex receives inputs from the dopamine cells in the tegmentum, as well as from the basal amygdala, ventral pallidum, and hippocampus. In addition, it sends outputs to the accumbens and to the motor cortex. It is thus in a position to receive information about behavioral arousal (dopamine connections from the tegmentum) and about conditioned

incentives and their amplification by dopamine (connections from the amygdala and pallidum). It can then integrate this information with data from long-term memory (connections with the hippocampus) and with the temporary contents of working memory (connections with other prefrontal areas) in the process of controlling movement (connections to the motor cortex). That the anterior cingulate is also involved in processing unconditioned incentives is suggested by studies showing activation of this region in humans receiving painful stimulation.

The orbital cortex, an area of the ventral prefrontal cortex, located at the bottom of the frontal lobe just above the eye sockets, is also important in motivation and decision-making. Studies of monkeys by David Gaffan and Edmund Rolls have implicated this region in emotional processing—in the processing of incentives (rewards and punishments) and in the temporary storage of incentive information. This region is connected with the anterior cingulate, and, like the anterior cingulate, it also receives information from the amygdala and hippocampus.

The work of Antonio Damasio has been especially significant in drawing attention to the critical nature of emotional information in human decision-making, and especially the role of the orbital cortex in this process. Much of Damasio's recent work on this topic is summarized in his two books, Descartes' Error and The Feeling of What Happens." He and his colleagues have shown that patients with damage to the orbital prefrontal region have poor judgment and often make decisions that lead to socially inappropriate courses of action. Using a clever experimental task, called the gambling task, they demonstrated that patients with orbital cortex lesions are insensitive to changes in the incentive value of stimuli, despite the fact that they have relatively normal working memory, as assessed in purely cognitive ways. Damasio and colleagues interpreted their results in terms of an inability of patients

with orbital cortex damage to use emotional information to guide action, and have proposed that emotional information or knowledge normally biases reasoning ability by influencing attention and working memory processes. Although they argued that their results dissociated decision-making from working memory processes, in light of the suggestion above that the orbital cortex is part of the working memory circuitry, and is especially involved in working memory about emotional information, an alternate conclusion is that the results show a dissociation between cognitive and emotional aspects of working memory.

The lateral, anterior cingulate, and orbital prefrontal regions are synaptically interconnected in various ways and should be thought of not as separate, independent modules, but as components of an integrated working memory system (chaps. 7 and 8). Although individual regions may make a relatively greater contribution to different aspects of working memory, especially in controlled laboratory experiments, in more natural settings, where decisions involve the integration of cognitive and emotional information, it is likely that interactions rather than dissociations between the

areas will be commonplace.

But we have to be careful not to jump to the conclusion that decision-making is simply something carried out by the prefrontal cortex on its own. Recent studies by Paul Glimcher at NYU have pointed to the contribution of a region in the parietal cortex that appears to be a component of the decision-making process by which eye movements are controlled. This region is part of the “where” pathway (chap. 7), is strongly connected with the lateral prefrontal cortex, and has long been implicated in eye-movement control. But Glimcher has upped the ante in this field by performing a sophisticated mathematical analysis that suggests that parietal neurons participate in decision-making. Basically, he’s approached the problem of how populations of neurons in the brain make decisions in a manner similar to the way that economists approach the behavior of populations of people. His analysis suggests that these cells are able to integrate information about a given stimulus and its implications for the amount of reward that a subject animal can expect on a particular trial, given what has been experienced in the past. It will be interesting to see whether this kind of analysis also works for cells in the prefrontal cortex and whether there are differences between parietal and prefrontal involvement. Also important to examine is whether there might be differences within prefrontal circuits. It might be expected that the medial or orbital areas would be more involved when emotional information is guiding decisions and the lateral prefrontal areas when emotional information is less important than cognitive information.

Many in the field have taken notice of Glimcher’s work. Particularly interesting is his claim that his approach can put Descartes’s dualism to rest—there’s no mental stuff living in the neurons and making decisions, it’s just the neurons carrying out mathematical calculations. While many have argued that a process of this type must underlie the mind-brain relation, Glimcher has actually eavesdropped on neurons while they do the calculations that we call decisions.

MOTIVATED MAN

Much of the work on motivation that we’ve considered here has been based on studies of animals. How far, you might ask, can this approach take us in understanding motivation in the daily life of human beings? The answer is, I think, fairly far. To understand why this is so, we need to consider two broad approaches to human motivation within psychology today.

The first, and most prominent, is the cognitive approach. In the 1950s, the social psychologist Leon Festinger fused the newly emerging cognitive view of the mind with the classic notion of drive in an attempt to account for the motivation of human behavior, especially in conflict situations. The theory, called cognitive dissonance, assumed that having inconsistent or conflicting thoughts (for example, “Wanting more money is greedy?” and “Having more money is good”) puts one in an uncomfortable state, a drive state, which demands reduction by changing one of the thoughts (for example, “Being greedy is not so bad” or “Having lots of money is not so good”). This theory gave rise to a tremendous amount of research in social psychology and spawned many other related theories.]

Social psychologists were one of the few groups of psychologists who remained interested in motivation during the cognitive revolution. But, over the years, their views lost their link to traditional notions of motivation and went farther and farther in the cognitive direction. Knowledge, beliefs, expectations, and self-awareness replaced drives and emotions in explanations of how behavior is aroused and directed.

For example, Nancy Cantor and Hazel Markus, leading contemporary cognitive motivational theorists, view motivation as a product of self-knowledge.’ Self-knowledge obviously includes knowledge about one’s emotions and their motivational consequences, but for Cantor and Markus motives are products of the self as much as contributors to it. Key to their theory is the notion of a working self, an on-the-fly construction about who we are that reflects who we’ve been (past selves), and who we want and don’t want to be (future selves). In contrast to earlier self theorists, such as Alfred Adler—who viewed the self as a static and enduring entity—Cantor and Markus view the self as a dynamic and mutable construction that changes in different situations—we have different goals when at home than

when on the job, for example, and the working self in each situation reflects such differences. One's working self is thus a subset of the universe of possible self-concepts that can occur at any one time—it is the subset that is available to the thinking conscious person at a particular moment, and is determined in part by memory and expectation, and in part by the immediate situation. These features of the working self explain how one can have both stable and mutable motives, and how motives can be conflicting or dissonant. The working self is a central part of one's mental apparatus. It influences perception, attention, thinking, memory retrieval, and storage, and guides action.

Building on William James, the cognitive psychologist John Kihlstrom has argued that a critical component of any conscious experience is the relation, in working memory, of the object of the experience to one's sense of self. Combining James and Kihlstrom with Cantor and Markus leads to the conclusion that the working self is constructed in working memory and the construction of the moment contributes in significant ways to on-line processing, decision-making, and behavioral control.

The second major approach to human motivation also originated in the 1950s when David McClelland proposed his need-achievement theory. According to McClelland, humans have a limited number of inborn motives (drives or needs) that are sensitive to natural incentives (hunger, thirst, and sex, for example), and motivation occurs when a person does something to make these goal objects available, usually by performing some behavior that has been learned in the past. When the behavior leads to the incentive, emotion is aroused, and the behavioral sequence that led to the incentive is reinforced. Through experience, people also learn to recognize cues that signal the availability of natural incentives. These cues also arouse emotions. Because positive emotional states are reinforcing, people seek out situations in which learned or natural incentives are present. Motives, in McClelland's view, are thus emotionally charged states that anticipate goal objects. Upon this quasi-biological foundation of inborn motives and learned incentives, McClelland has built a theory of human behavior focused on positive motivational states like affiliation and achievement. McClelland does not deny the importance of negative motivation, but places this topic outside the area covered by his theory. Research on humans by McClelland and others accumulated over several decades has tended to support major aspects of achievement theory.

So how do we reconcile the two views of human motivation? McClelland's theory emphasizes nonverbal motivational systems, some of which are biologically organized. These work implicitly and function more or less similarly in humans and other mammals, though each species can have its own specialized motive systems as well. McClelland's views fit well with the drive/incentive/reinforcement tradition. But this approach does not cover the full

range of human motivation, as McClelland himself has noted. Consciously accessible, verbally encoded, explicit motives acquired through the use of language are also an important part of human mental life, especially in social situations. The self-consciousness view of motivation thus complements rather than contradicts the implicit view.

Anthropologists Claudia Strauss and Naomi Quinn have used McClelland's rapprochement between implicit biological and explicit psychological motives to account for motivational similarities and differences between cultures. She points out that some motives are universal and vary only in intensity between cultures, whereas others are unique to specific cultures. One obvious way to reconcile these theories is to place McClelland's biologically based motives in the former category and the explicit, self-conscious, individually constructed motives of Cantor and Markus into the latter. However, not all learned motives are related to explicit knowledge. Most motivational theories we've discussed, from Thorndike to Hull and Miller to McClelland, emphasize the learning of motivational habits. Habit learning generally is viewed as an implicit form of learning. Strauss points out that habits are learned not just through encounters with reinforcers (primary or secondary) but also in common social situations in which we observe the successes and failures of others. Indeed, in spite of the past emphasis of social psychology on self-conscious motives, more and more research on social behavior is beginning to emphasize implicit or unconscious aspects of human motivation. Particularly noteworthy is John Bargh's work on *automotives* and Tim Wilson's theory of hidden selves. So it also needs to be pointed out, though, that while habit systems may indeed learn implicitly, when motivational habits are learned under conditions in which working memory attends to the present circumstances, the

consequences of habit learning can be explicitly represented and consciously knowable, and only later, once routinized, sent to the depths of the mind. Because habits can be both useful and pathological (habitual avoidance of the outside world by a patient with panic disorder, for example), they are an especially important form of behavioral learning.

Both implicit and explicit systems thus contribute to motivation. Working memory is important in the guidance of behavior toward goals that are explicitly represented there and can be carried out under executive control functions. But, at the same time, we also have brain systems that work implicitly in the processing of incentives and the guidance of behavior toward goals. Sometimes, implicit and explicit motivation will be in sync, in which case working memory and implicit systems guide behavior with a common put pose. For example, following the explosion of a bomb, it is likely that both explicit and implicit systems will initiate behaviors that take you away from rather than toward the source of the explosion, all other things being equal. But other things are not always equal. If your spouse had gone to the concession stand to get a snack, and the stand is in the direction of the explosion, you may, through an executive decision, override the tendency to flee and instead run toward the explosion. Indeed, functional-imaging studies by Jonathan Cohen suggest that the prefrontal cortex, especially the anterior cingulate, is involved in resolving motivational conflict. This region and its synaptic connections may therefore play a crucial role in overcoming fear or other emotional states when we need to take an action that goes against our innate or learned compulsions. (Might this circuitry also underlie cognitive dissonance?)

An informed conception about the brain regions and circuits underlying a particular function is the best way to begin to determine how the brain actually performs that function. Therefore, once we've translated basic motivational concepts into plausible neural circuits, the task of relating the enormously complex topic of human motivation to the brain becomes less daunting.

THE MENTAL TRILOGY IN ACTION

Through the topic of motivation, we begin to see the mental trilogy in action. A mind is not, as cognitive science has traditionally suggested, just a thinking device. It's an integrated system that includes, in the broadest possible terms, synaptic networks devoted to cognitive, emotional, and motivational functions. More important, it involves interactions between networks involved in different aspects of mental life.

Often the things we attend to and remember are the things that are important to us. In such situations, cognitive processing will be accompanied by emotional arousal. And emotional arousal does not stop with a simple reaction, for we often use it to guide our behavior toward or away from the situation that the emotionally arousing stimulus signifies. In the process of doing so, we sometimes have to make decisions about what to do in order to keep working toward our goal. The goal has to be kept in mind, even when detours have to be taken. But the system also needs to be reset if a more important goal emerges along the way. This requires reallocation of cognitive, emotional, and motivational resources, and adjustments within and between the component processing systems.

Self-knowledge is certainly a significant aspect of human motivation, but even animals that are not self-aware, or at least not robustly aware of who they are the way a human is, are motivated to do things—they seek food and shelter and avoid predators and injury. Much of what we humans do is also influenced by processes that percolate along outside of awareness. Consciousness is important, but so are the underlying cognitive, emotional, and motivational processes that work unconsciously.

CHAPTER TEN

SYNAPTIC SICKNESS

CANST THOU NOT MINISTER TO A MIND DISEASED,
PLUCK FROM THE MEMORY A ROOTED SORROW,
RAZE OUT THE WRITTEN TROUBLES OF THE BRAIN,
AND WITH SOME SWEET OBLIVIOUS ANTIDOTE
CLEANSE THE STUFF'D BOSOM OF THAT PERILOUS STUFF
WHICH WEIGHS UPON THE HEARP

—Shakespeare, *Macbeth*.

In spite of Shakespeare's insight that mental problems are "troubles of the brain," hundreds of years later, proponents of a physical basis for mental illness still had little evidence with which to make their case. By the late nineteenth century, it was widely accepted that actual destruction of the brain (the result of diseases like syphilis) could change mental function drastically,⁷ but the alterations of mood and thought that had come to be known as neuroses and psychoses were still resistant to physical explanation. Young Sigmund Freud sought a neurological account of mental illness, but realizing that such a goal was unattainable in his lifetime, turned to psychological explanations and psychological treatments instead. There was no "sweet oblivious antidote," to use Shakespeare's phrase, for hysteria, melancholia, or anxiety.

Obviously, the situation has changed considerably in the past few decades. Mental maladies have, in the eyes of many, come to be recognized as "troubles of the brain," and antidotes have emerged as the treatment of choice more often than not. Regardless of how one feels about the biological orientation of psychiatry today (and it has many critics) two facts must be acknowledged. The essence of who we are is encoded in our brains, and brain changes account for the alterations of thought, mood, and behavior that occur in mental illness. The key issue is not whether mental illness is really neural in nature. It is instead the nature of the neural changes that underlie mental problems, and the manner in which treatment should proceed.

Biological psychiatry was founded on, and still largely adheres to, the assumption that mental disorders are due to chemical imbalances in the brain. The brain, in this view, is like a delicate soup. To maintain its characteristic flavor, just the right mix of key ingredients is required. Too little or too much of one or another item, and its character is altered. But the normal character can sometimes be restored with the right mix of additional ingredients that either replace a missing ingredient or dampen the effects of an overly abundant one. Like a master chef's, the biological psychiatrist's job is to adjust the blend, to balance the chemistry, so that the desired character is restored.

The soup model evolved in psychiatry for two reasons. One was practical—a number of mental disorders were found, sometimes accidentally, to be helped by altering brain chemistry. Another was conceptual—neuroscientists had come to think of chemicals as an important contributor to the way in which mental life is created by the brain. The practical justification remains, but, as we'll see now, the conceptual one has changed.

Although it seemed in the early 1960s that there might be a particular chemical code for different mental

states —molecules of pleasure, fear, and aggression were proposed—this notion is now viewed by many as naive. Mental states are not represented by molecules alone, or even by a mix of molecules. As we've seen, they are instead accounted for by intricate patterns of information processing within and between synaptically connected neural circuits. Chemicals participate in synaptic transmission, and in the regulation or modulation of transmission, but it is the pattern of transmission in circuits, more than the particular chemicals involved, that determines the mental state. The old battle cry "No twisted thought without a twisted molecule needs replacing. Synaptic changes, not molecules, underlie mental illness. The importance of circuits has been illustrated repeatedly in the previous chapters. In this chapter, I want to emphasize the implications of a circuit point of view for understanding the nature of psychiatric disorders.

Biological psychiatrists do realize that circuits are significant, and it's not out of ignorance that they adhere to the soup model. They would love to have smart drugs, akin to smart missiles, that, when taken orally, would go straight for the circuits that underlie a particular disorder and avoid all others. With this sort of drug, dreaded side effects would be eliminated. But achieving such precision would require that the exact circuits altered in particular mental disorders be known, and that drugs be available to adjust selectively the chemistry of the relevant circuits. Although not all of the pieces of the puzzle are in place today, research efforts are under way to obtain the information that would make this possible. As Steve Hyman, director of the National Institute of Mental Health, has noted, psychiatry, arm in arm with neuroscience, is poised to answer many of its central questions. Hyman, in fact, refers to the time ahead as "the millennium of mind, brain and behavior." To facilitate the achievement of his goal, the National Institute of Mental Health recently has established several Research Centers for the Neuroscience of Mental Disorders (I serve as the director of one of these, the Center for the Neuroscience of Fear and Anxiety, which I'll describe later in the chapter).

Most critics of biological psychiatry will not be satisfied by the shift in paradigm from soup to circuits. They will argue that the important question is not whether we should focus on connections instead of chemicals, but whether we should focus on the brain at all in our efforts to understand and treat mental illness. Psychological problems, such critics insist, are rooted in life experiences rather than in malfunctioning brain hardware, and should be treated by helping the patient come to terms with the underlying problem rather than by trying to alter brain circuits. But as the arguments presented in this book make clear, life's experiences leave lasting effects on us only by being stored as memories in synaptic circuits. Because therapy is itself a learning experience, it, too, involves changes in synaptic connections. Brain circuits and psychological experiences are not different things, but rather, different ways of describing the same thing. Still, the manner in which brain changes are effected by psychotherapy is not necessarily the same as the way in which they are by drugs, which is why therapy works in some cases where drugs don't and vice versa, or why a combination of drugs and psychotherapy can sometimes be more productive than either alone.

Enough generalities—let's examine what is known about the biology of mental illness. I can't cover all categories of mental illness here, and so will focus on schizophrenia, depression, and anxiety disorders: In assembling a coherent account of the history of drug therapy for these disorders, I've borrowed from books by Samuel Barondes (*Molecules and Mental Illness*) and Elliot Valenstein (*Blaming the Brain*).

TRIPPIN' UTERUSES

The mid-twentieth century was an exciting time in neuroscience research. With the end of World War II, national scientific resources were shifted to peacetime efforts, and one area that boomed was brain studies. The neuron doctrine had long since prevailed, and synaptic transmission between neurons was widely accepted as the currency of the brain. But the nature of synaptic transmission was still the topic of much investigation, the interpretation of which was hotly debated. Sir John Eccles pushed the notion of electrical transmission, while Sir Henry Dale emphasized the chemical nature of synaptic communication. In the end, both were right—some synapses are chemical and others electrical. Nevertheless, chemical transmission is believed to be the dominant form of synaptic

communication, a finding to which even Eccles came around. Still, it is a long leap from the idea that tiny amounts of neurotransmitter are squirted across individual synapses to the notion that mental states and mental disorders are chemically coded. Another set of events, culminating around the same time, played an important role in closing the gap.

Ingestion of a variety of substances from nature, like peyote or certain mushrooms, had long been known to induce hallucinations, and by 1937, it had been proposed that the hallucinations induced by peyote resembled the symptoms of schizophrenia. But it was the accidental discovery of a novel hallucinogen, d-lysergic acid, or LSD, that helped jump-start the psychopharmacological industry in the 1950s."

Chemists had determined in the 1930s that ergonovine was the active ingredient in the fungus ergot, which had been used by midwives to induce uterine contractions and control bleeding during childbirth. In the process of trying to formulate a new drug that would be useful for obstetricians, Albert Hoffman was accidentally exposed to one of the chemical compounds with which he was working and entered a "dream-like state" in which he experienced a "kaleidoscopic play of colors." Suspecting that the culprit was LSD (the twenty-fifth modification of an acidic extract of ergonovine), he decided to take a small amount to see if the experience would recur. Again, he had powerful hallucinations. Furniture in the room "assumed grotesque, threatening forms," and the woman next door became a "malevolent, insidious witch with a coloured mask." His feeling that a demon had invaded him lasted for about fourteen hours. Others at the company tried the substance as well, all with the same result. News got out, and so did the drug. The key event for our story was not the popularization of LSD by figures like Ken Kesey and Timothy Leary, but rather the discovery in 1953 by John Gaddum, a Scottish scientist, that LSD blocked uterine contractions induced by the chemical s-hydroxytryptamine, otherwise known as serotonin. Serotonin subsequently was found in the brain, and shown to be a neurotransmitter. These facts gave Gaddum an idea. If LSD produced psychotic states (hallucinations) by altering neurotransmission in the brain, sanity might require a certain level of brain neurotransmitters, and changing transmitter levels might be a way of treating mental illness.

That drugs might be useful in treating mental illness was not a novel idea. What was new was the notion that abnormal mental function might result from alterations in synaptic transmission, and that drugs that affected synaptic transmission might therefore be used to treat mental illness. Gaddum's hypothesis thus marks the beginning of the modern approach to biological psychiatry.

FROM SEDATION TO CORRECTION IN PSYCHOSIS

As many psychiatrists are aware, the neo-Freudian Harry Stack Sullivan encouraged his patients to partake of alcohol to loosen them up before starting psychoanalysis. This was not drug therapy per se, as it was not intended as a cure. In fact, during the first half of the twentieth century, the noncurative

use of drugs was common, not so much to facilitate therapy as to sedate agitated or manic patients and make them more manageable, especially in institutional settings.

Insulin was one of the first drugs conceived of as helping to cure a mental disorder. Like many other drugs, its therapeutic effects were discovered by chance. In 1933, a Viennese physician named Manfred Sakel gave a small dose of insulin, which lowers blood sugar levels and thereby stimulates appetite, to some of his agitated schizophrenic patients who had not been eating. Surprisingly, their psychotic symptoms were somewhat relieved. When Sakel tried a larger dose, the extreme lowering of blood sugar put the patients in a coma for several days, but when they recovered, their psychotic symptoms were much improved. Although the reason this treatment worked was a complete mystery, insulin coma therapy came to be widely used, in spite of its dangers, since there was no other viable treatment for schizophrenia.

In the 1950s, though, two new psychotic treatments emerged, both of which were explained in terms of new discoveries about neurotransmission. One of these was reserpine, a drug isolated from a plant called *Rauwolfia serpentina*. *Rauwolfia* had been used to treat insomnia and insanity in India since ancient Hindu times. When Indian doctors in the 1930s found that *Rauwolfia* could reduce high blood pressure, drug companies were motivated to

isolate reserpine, the active substance in the plant. In the meantime, Nathan Kline, a New York psychiatrist who knew of the use of *Rauwolfia* to treat insanity in India, obtained reserpine and administered it to schizophrenic patients. He noted that the drug made the patients less suspicious and more cooperative. Soon thereafter, Kline convinced the governor of New York that reserpine should be given to the 94,000 patients in the massive New York State Psychiatric Hospital system. Klein's report came right around the time that Gaddum had made the connection between serotonin and sanity, and it wasn't long afterward that scientists determined that reserpine reduced the level of monoamines (including serotonin, norepinephrine, and dopamine) in the brains of experimental animals. By the mid-1950s, it seemed that psychosis was indeed related to alterations in monoamine levels in the brain, and that restoration of monoamine function might therefore be a way of treating this condition.

This conclusion was supported and refined by the near-simultaneous discovery of the beneficial effects of chlorpromazine, a phenothiazine drug marketed as Thorazine. This drug was isolated because its chemical structure was similar to that of antihistamines. A French surgeon tried administering chlorpromazine to some of his patients to test whether it might be more effective than other antihistamines in sedating patients and reducing respiratory complications in surgery. The drug turned out to have a powerful calming or tranquilizing effect, which led the surgeon to propose using it with psychotics. Two French psychiatrists, Jean Delay and Pierre Deniker, tried administering chlorpromazine to a few schizophrenic patients who were resistant to all other forms of treatment. Not only were they tranquilized, their psychotic symptoms (paranoia and hallucinations) were also reduced.

One of the side effects of both reserpine and chlorpromazine, however, was the development of a movement disorder called tardive dyskinesia, involving muscle rigidity similar to that typical of Parkinson's disease. Since Parkinsonism had been found to be associated with a reduction in dopamine in the brain, the suggestion that antipsychotic drugs were useful in treating schizophrenia because they altered dopamine transmission arose, as did the related notion that schizophrenia was due to excess dopamine. In fact, it came to be believed that the potency of antipsychotic drugs was directly related to their ability to produce tardive dyskinesia. The view that schizophrenia involved dopamine transmission was strengthened by the complementary observation that psychotic symptoms are brought on by too high a dose of L-dopa, which increases dopamine levels and is used to treat Parkinson's patients, or by amphetamine (speed), which artificially stimulates dopamine receptors.

An experiment performed by a Swedish scientist, Arvid Carlsson, helped to relate the effects of reserpine and chlorpromazine to dopamine. He examined the effects of each drug on the level of the various monoamines in rats. While reserpine caused a reduction, chlorpromazine had no effect. This was a troubling finding for the emerging dopamine theory of schizophrenia. But in a follow-up experiment, he found that the breakdown products of dopamine, and not the other monoamines, were increased. This implied that more dopamine was being released, possibly to compensate for the block of dopamine receptors by chlorpromazine. The logic is complex, but the important point here is that this study suggested that both antipsychotic drugs acted by reducing dopamine transmission in the brain. While reserpine worked by reducing dopamine release from presynaptic terminals, chlorpromazine's effect was due to a blockade of postsynaptic dopamine receptors. Carlsson shared the Nobel Prize in 2000 for this and related work.

The idea that schizophrenia might be explained at the level of dopamine receptors was strongly supported by research carried out by the Dutch scientist J. M. van Rossum. Rats given amphetamine displayed increased locomotor activity. Van Rossum showed that both the antipsychotic drugs chlorpromazine and reserpine blocked this increased activity. Although these two drugs affect all monoamines, a third potent antipsychotic drug, haloperidol (Haldol), which only affected dopamine, was particularly powerful in reducing the locomotion. Consequently, van Rossum proposed that anti-psychotic drugs do their work by blocking dopamine receptors, and that overstimulation of dopamine receptors was at the heart of schizophrenia.

Still, the evidence linking dopamine neurotransmission to schizophrenia was circumstantial. An additional discovery strengthened the connection considerably. By the mid-1970s, a number of antipsychotic drugs had been developed. Studies by Philip Seeman in Toronto and by Ian Creese and Solomon Snyder in Baltimore rank-ordered the drugs by their clinical effectiveness and then examined the ability of each to block dopamine receptors.⁹ Both

studies found that the correlation was nearly perfect: the more potent its dopamine-receptor-blocking ability, the more effective the drug was in treating schizophrenia. The demonstration that the same relation did not hold between clinical effectiveness and blockade of serotonin or norepinephrine bolstered the view that it was dopamine, and not the other monoamines, that was involved]" Dopamine receptors come in two groups: D₁-like and D₂-like. The fact that all of the classic dopamine-receptor-blocking drugs were found to produce their antipsychotic effects by acting on the D₂ class added further specificity to the role of dopamine in schizophrenia]"

The bottom-line conclusion from the early work on the relation of monoamines to schizophrenia was this: Too much dopamine induces a psychological disorder, and too little, a movement disorder similar to Parkinson's disease. A balance must be maintained. When it is tipped, the brain does not function normally. Subsequent work, as we'll see in the following section, has drawn a more complex picture of the neural basis of schizophrenia. Dopamine is still believed to be involved, but not quite in so simple a way as the original imbalance hypothesis suggested.

RETHINKING PSYCHOSIS

Although blocking the action of dopamine at its receptors in the brain is an effective way to treat psychotic symptoms, and although artificial stimulation of dopamine receptors (with L-dopa, amphetamine, or cocaine) can bring on psychotic symptoms, the dopamine receptor hypothesis of psychosis ran into trouble for several reasons. First, the therapeutic effects of dopamine drugs are far too slow to be accounted for by a simple receptor process. The drugs bind to receptors as soon as they are released from the presynaptic terminal, and have their physiological effects moments thereafter. But actual therapeutic effects of these substances take a week or two (or longer) to begin. The therapeutic effect is therefore more like a long-term adjustment to the receptor blockade than a direct and immediate effect of the blockade. Second, attempts to demonstrate that dopamine levels are elevated in schizophrenia proved unfruitful, in contrast to the relative ease with which decreases in dopamine were found in Parkinson's disease. The search for changes in dopamine receptors, rather than an increase in dopamine itself, was more successful² but the evidence was controversial, since most of the research was performed on patients who had been on antipsychotic medication, which makes it difficult to know whether the change in receptors was due to the disorder itself or was a consequence of the treatment]. Third, standard manifestations of schizophrenia include positive or added symptoms (hallucinations, abnormal thought patterns, paranoia, delusions, agitation, hostility; bizarre out-of-context behavior) and negative symptoms or deficiencies (blunted emotions, cognitive deficits in attention and working memory, poor hygiene, poverty of speech, social isolation, loss of motivation). The classic or typical antipsychotic medications, D₂ receptor blockers, are mainly useful in treating positive symptoms, and are less helpful in treating negative ones, suggesting that the dopamine theory is at best incomplete. Finally, recent studies have shown that schizophrenia can also be helped by drugs that target other systems, including serotonin and norepinephrine. These so-called atypical anti-psychotic medications turn out to be more useful for negative symptoms, further challenging the sufficiency of the original dopamine theory.

Recent efforts to explain schizophrenia and other mental disorders have turned to more complex conceptions that focus on alterations of function in specific brain regions and circuits rather than global changes in the level of monoamines. For example, a number of studies have demonstrated structural differences between the brains of schizophrenics and normal control subjects. Included are changes in size or volume of certain brain regions, and the number of cells and their shape and arrangement in different regions. Some of the key areas in which structural changes have been discovered include the prefrontal cortex and medial temporal lobe (hippocampus and amygdala). Changes in the number of dopamine receptors have also been found in the prefrontal cortex and basal ganglia. The prefrontal cortex, hippocampus, and amygdala also exhibit functional anomalies in blood flow

and/or neural activity, as determined by PET and functional MRI scanning techniques]

The structural and functional differences in the prefrontal cortex have attracted considerable attention and have led to a revised dopamine theory of schizophrenia. Classic or typical antipsychotic drugs are believed to achieve their beneficial effects on positive symptoms by blocking D₂ receptors in the basal ganglia. Indeed, a number of

studies have found that D2 receptors are elevated in the basal ganglia of schizophrenics,³⁰ including patients not previously medicated.] However, as we have seen, D2 blockade has little effect on negative symptoms. Some of the most prominent negative symptoms include cognitive changes, and the degree of cognitive deficit is directly related to treatment prognosis. It has been proposed that deficits in working memory underlie the cognitive changes in schizophrenia. Working memory, you'll recall, crucially involves the prefrontal cortex. Given that there are relatively few D2 receptors in the prefrontal cortex, it is not surprising that D2-related drugs have little effect on cognitive negative symptoms.

At the same time, the D₁ class of receptors is abundant in the prefrontal cortex of normal persons but is reduced in schizophrenics. Further, working memory performance in schizophrenics is related to the number of D₁ receptors in the prefrontal cortex—the fewer the number of receptors, the poorer a patient performs. Moreover, cognitive tests that require temporary storage and executive function typically result in increased functional activity in the prefrontal cortex in normal humans, but not in schizophrenics. Consistent with these results is the fact that blockade of D₁, but not D₂, receptors in the prefrontal cortex of rhesus monkeys impairs delayed response performance (a cognitive test of the temporary storage aspect of working memory).

The revised version of the dopamine theory thus proposes that schizophrenia involves overactivity of the D₂ class of receptors in the basal ganglia and underactivity of D₁ receptors in the prefrontal cortex. The hyperactivity of D₂ receptors in the basal ganglia accounts for positive symptoms, and the underactivity of D₁ receptors in the prefrontal cortex accounts for negative symptoms. This explains why drugs that block D₂ receptors only treat positive symptoms, but leaves open the question of how atypical antipsychotic drugs, which mainly target serotonin or norepinephrine systems, treat negative symptoms.

The improvement in negative symptoms is believed to come about because the atypical medications cause an elevation in the amount of dopamine released from terminals in the prefrontal cortex and thereby help overcome the deficits caused by D₁ understimulation. The site of the drugs' action is not in the prefrontal cortex itself, but in the ventral tegmental area, which is the home of the neurons whose terminals release dopamine in the prefrontal cortex. For example, some very effective atypical antipsychotic drugs block the serotonin 2A receptors, which are located on ventral tegmental dopamine neurons. When the serotonin receptors are stimulated, they inhibit the ability of the tegmental cells to release dopamine in the prefrontal cortex, and other areas. By blocking serotonin receptors, atypical antipsychotic drugs block the inhibition normally produced by serotonin. This leads to an increased firing of the tegmental cells and thus more dopamine release in the cortex. The extra dopamine released helps compensate for the reduced number of D₁ receptors in the schizophrenic brain.

Research on antipsychotic drugs is currently advancing in many directions. In addition to work on drugs that target monoamines, much work is attempting to determine whether alterations in glutamate and/or its receptors, especially NMDA receptors, may help treat schizophrenia. Other studies are examining the role of GABA transmission.

In general, there is a growing interest in the idea that alterations in synaptic connectivity in neural circuits, rather than just levels of neurotransmitters or receptors, are important. Neurotransmitters and receptors still figure prominently in this approach, but in the context of connections within and between areas. For example, Francine Benes, noting that the number of GABA cells is decreased in the hippocampus of schizophrenics, has proposed that the disorder involves a shift of dopamine connections from excitatory glutamate cells to GABA cells in the hippocampus to compensate for a loss of inhibitory cells, and David Lewis has proposed that a specific class of GABA cells located in a specific layer of the prefrontal cortex plays a critical role in cognitive regulation and is altered in schizophrenia. In imaging studies, researchers such as David Silbersweig and Emily Stern are beginning to examine changes in neural activity in networks involving interconnections of the prefrontal cortex with other cortical and subcortical areas.

Perhaps the most sophisticated circuit theory of schizophrenia to date has been the work of Anthony Grace. Based on his studies of dopamine in various forebrain circuits, he has proposed a complex theory that attempts to ex-

plain how synaptic interactions between the prefrontal cortex, hippocampus, and amygdala account for certain aspects of schizophrenia. Recall from our discussion of motivation circuits in chapter 9 that the prefrontal cortex sends connections to the nucleus accumbens, which connects back with prefrontal

areas by way of the ventral pallidum. This is potentially a significant set of connections, since it includes the key areas involved in both negative and positive symptoms. The prefrontal cortex, as we've seen, has abundant D1 receptors that seem to be involved in negative symptoms. Similarly, the nucleus accumbens is one of the key regions of the basal ganglia that receive dopamine inputs from the ventral tegmental area and is loaded with D2 receptors, which, as we've also seen, are involved in positive symptoms.

Grace showed that the ability of the prefrontal cortex to activate the accumbens was regulated by the hippocampus and amygdala. That is, when there was synaptic activity from the hippocampus or the amygdala to the accumbens, the prefrontal cortex could fire accumbens cells; otherwise, it could not. The hippocampus and amygdala thus gate or regulate the prefrontal activation of accumbens. Given the role of the hippocampus in contextual processing of the environment, and the role of the amygdala in emotional processing, Grace proposed that these connections allow prefrontal activity to be adjusted in response to changes in the global environmental situation as well as to specific emotional stimuli. In schizophrenia, alterations in the hippocampus and amygdala (both of which have been found) could affect prefrontal processing in the accumbens, reducing the ability of the patient to produce responses appropriate to the immediate cognitive and emotional situation.

In Grace's theory, dopamine modulates the synaptic interactions in all of the regions. Specifically, Grace has proposed that there are two different pools of dopamine, which have different roles. The most obvious pool is the dopamine that is released by axon terminals in the striatum or prefrontal cortex when dopamine cells in the brain stem fire action potentials. This is called the phasic pool of dopamine. Some of the dopamine released phasically does not get broken down or sucked back up but instead diffuses away from the immediate area of the synapses and collects over time. This chronic pool of dopamine will affect dopamine receptors that are in its vicinity even if they are not at the moment receiving phasic dopamine from the brain stem cells. Grace argues that the chronic dopamine is what is altered over the course of long-term drug treatment, which is why it takes so long for such treatment to be effective. These alterations, in turn, adjust the phasic actions of dopamine and influence processing in the circuits just described.

Clearly, research in this area has made much progress. Although the situation has grown more complicated now that schizophrenia is no longer thought of as simply a problem of too much dopamine, the puzzle is more likely to be solved by accepting complexities than by ignoring them. Circuit models, for example, are more consistent with the fact that schizophrenia is not a single disorder but a family of conditions that can be manifested in different ways in different people. The particular circuits affected, including transmitter alterations, may well determine the extent to which positive and negative symptoms appear and the manner in which they will be expressed. Key to future breakthroughs may be the ability to relate specific symptoms to specific circuits, and the ability to develop drugs that target those circuits.

BALANCING DEPRESSION

Around the same time that drugs were being developed to treat schizophrenia in the 1950s, depression, too, was proving susceptible to chemical ameliora

tion. Although it was not realized at first, monoamines were also involved in this illness.

One of the first drugs successfully used in treating depression was iproniazid. This compound was developed from hydrazine, which was originally used as rocket fuel by the Germans in World War II. After the war, other compounds based on hydrazine were developed and tested for medical applications. One of these, iproniazid, was found to inhibit the bacteria that give rise to tuberculosis. Then TB patients were treated with the drug, they were also observed to become euphoric. On the basis of this effect, iproniazid was given to depressed patients. Although the results were not convincing, Nathan Kline, the man who gave reserpine to schizophrenics in the New York State

system, decided to give iproniazid a try. He put a diverse group of patients on the drug, and depressed ones improved. The key difference between Kline's study and the earlier one was that he kept the patients on the drug for five weeks; it is now commonly accepted that antidepressant drugs take some time to kick in, so Kline's decision to extend treatment was crucial. According to Kline, within a year of reporting his results, 400,000 depressed patients had been treated with iproniazid.

The mechanism by which iproniazid worked was unknown, but studies of animal brains soon revealed that monoamines were involved. Normally, the enzyme monoamine oxidase causes monoamines to be degraded after they are released from axon terminals, helping to end the action of the monoamines at postsynaptic receptors (stopping transmission is as important as starting it). Iproniazid was found to prevent the breakdown of monoamines after they are released. As a result, the monoamines remain around the synapse longer, and have a greater effect. Because iproniazid achieves its effects by inhibiting monoamine oxidase, it is called a monoamine oxidase inhibitor (MAO inhibitor).

The MAO inhibitors quickly became a common form of treatment for depression. But the treatment, successful or not, had a cost. MAO is a player in many brain areas, not just those that are altered in depression, and also participates in various functions outside the brain. When a drug is taken orally, it affects all relevant functions and not just those the patient wants corrected. One of the roles of MAO is the breakdown of the amino acid tyramine, which enters the body from ingested food substances (for example, certain cheeses). Unchecked, a rise in tyramine in the body can lead to a sudden and life-threatening rise in blood pressure.

The other major treatment for depression that emerged in the 1950s involved a class of drugs known as tricyclics (due to the three rings of their chemical structure). One of these substances, imipramine, was initially tried as a treatment for schizophrenia, due to some similarity to phenothiazines. However, because the drug seemed not to calm the patients but to energize them and elevate their mood, it was tested as a treatment for depression, and positive results were obtained. This class of drugs had fewer side effects than the MAO inhibitors, and tricyclics with brand names like Tofranil and Elavil became widely used to treat depression.

Like the MAO inhibitors, the tricyclics work by maintaining elevated levels of monoamines at postsynaptic sites. But the mechanism of their activity is different. The action of monoamines, once released, is terminated not just by the substances' being broken down, but also by being transported from the synaptic space back into the presynaptic terminal. Tricyclics block the transportation step, thereby allowing the monoamines to remain available to bind to postsynaptic receptors for longer periods of time. Their effect was discovered by Julius Axelrod, who received a Nobel Prize for this and other work on a monoamine chemistry. Because the side effects of tricyclics were less severe than those of MAO inhibitors, they soon came to be the first line of attack against depression.

By the mid-1970s, it seemed that depression was due to too little monoamine transmission, just as schizophrenia was due to too much. This conclusion implied a rather simple, in fact simplistic, picture of the illness. We now know that its mechanism is much more complicated. When considered from the point of view of whether patients were helped, rather than whether the underlying theory was correct, the monoamine hypothesis was a success. But the success came not from the fact that the monoamine drugs that were developed in the 1950s, 1960s, and 1970s were the ultimate answer. Instead it was because they helped pave the way for a logical approach for future drug development—by making effective drugs more specific, side effects might be reduced. As in the case of schizophrenia treatment, recent attempts to develop drugs for depression are more promiscuous in their flirtations with transmitter systems in an attempt not just to reduce side effects but to create more effective treatments. This has severely compromised the monoamine theory, but has facilitated the search for new and more effective approaches to treating depression.

DRUGS THAT TALK

In the late 1980s, a new chapter in the treatment of depression began with the development of drugs that selectively enhanced the availability of serotonin,

in contrast to the earlier generation of drugs, which augmented both serotonin and norepinephrine. Offering the relief of depression with fewer side

effects, and possibly even making “normal” people happier, the selective serotonin reuptake inhibitors (SSRIs) took the world by storm. Cover stories in magazines like *Newsweek* and a best-selling book, Peter Kramer’s *Listening to Prozac*, helped make them a popular phenomenon. SSRIs have been prescribed to millions and millions of people around the world, not just for debilitating depression but also for various shades of unhappiness, back pain, anxiety, stress, PMS, and a variety of other discontents. Prozac came to be seen as the ticket to a happier life, free of stress and filled with friends and success. As *Newsweek* put it, “Shy? Forgetful? Anxious? Fearful? Obsessed? How Science Will Let You Change Your Personality with a Pill.”

Over the past decade or so, a number of different SSRIs have been marketed, including Prozac, Zoloft, Paxil, Luvox, and Celexa. These drugs are as effective as tricyclics in helping depressed people, but probably not more so. This is to be expected, since both classes of drug basically do the same thing—make more serotonin available at synapses. But the SSRIs do the job with fewer side effects—they are selective for serotonin, so the side effects caused by enhancing norepinephrine are eliminated.

More people have taken SSRIs than any other prescription medication. About 3.5 million prescriptions were written in one year alone. Because the side effects are less severe than those of MAO inhibitors or tricyclics, patients don’t have to be monitored so closely, and many prescriptions are given out by general practitioners rather than psychiatrists.

There is, however, a growing movement against the use of SSRIs and other psychiatric medications. In books like *Talking Back to Prozac* and *Toxic Psychiatry*, Peter Breggin has mounted an attack against the widespread use of drugs, arguing for an approach to psychiatry based on love, trust, understanding, and traditional psychotherapy.” His Website claims that the side effects of SSRIs include violent and suicidal tendencies, and cites all sorts of literature on the topic, including articles with titles like “Prozac and Xanax Found by Court to Cause Criminal Conduct” and “Was School Shooter Eric Harris Taking Luvox?” Side effects of SSRIs, including facial tics and sexual dysfunction, are also emphasized in *Prozac Backlash*, by Joseph Glenmullen. While these are legitimate concerns, it is important that they be evaluated in the same way as the positive effects of the drugs, that is, through controlled studies rather than on the basis of anecdotes.

Yet another group of books makes a personal case for the use of such drugs. It’s hard not to be moved by the autobiographical accounts of people who, after years of suffering with depression, have been given a new lease on life by drug therapy. Kay Redfield Jamison’s *An Unquiet Mind*” is especially noteworthy, but there are others. Again, though, careful research is the best way to evaluate drug efficacy.

What, then, are the chances that these drugs actually will help an individual? Scientific studies show that many clinically depressed patients (about 50 to 60 percent) improve when given SSRIs or tricyclics, but some people given placebos (sometimes as many as 30 percent of the sample) get better as well. The most cynical view of such data is that drugs don’t accomplish much more than placebos. The more positive view—that some people really are helped by drugs—is suggested by the fact that in contrast to SSRIs, placebo effects tend to be smaller and shorter-lasting and hard to interpret. Further, a recent report has challenged the validity of placebo effects.

But what about the notion, promoted by critics, that many of the people taking SSRIs are not really clinically depressed but just out for a little mood boost or personality makeover? I asked Greg Sullivan, a psychiatrist friend of mine from Columbia University; what he thought of this criticism, and he responded as follows:

In my experience I do not know of anyone for whom this is true. In general, these drugs are not fun to take (it’s bothersome to remember a daily pill, they often cause sexual dysfunction, they can lead to slow weight gain, and friends and family sometimes imply that taking drugs for depression is a sign of mental weakness). If someone is pleased with the effects of an SSRI, that’s usually an indicator that the drug has had a significant impact on serious

symptoms, including those caused by a chronic low-level depression (dysthymia). Such people are amazed once they learn that low energy, chronic sleep disturbance, and a pessimistic approach to just about everything doesn't have to be the norm for the rest of their life. But SSRIs are not "happy pills," and people without significant mood or anxiety disturbance do not generally get anything beneficial from them, certainly nothing that would make them sustain the use. This is a sharp contrast to psychiatric drugs that are sometimes abused such as Ritalin, which makes people high if they administer enough, or benzodiazepines, which some people use in order to become numb to their difficulties and to life in general.

The bottom line seems to be that, if you are seriously depressed, drug treatment, especially in conjunction with psychotherapy, is probably your

best hope, especially if psychotherapy alone has been unhelpful. There are likely to be side effects, and you have to decide whether the side effects are going to be worse than living with depression. But this is a question that has to be posed about almost any medication, since there are few magic bullets in medicine.

By now, SSRIs have been in the spotlight for over a decade and have been very profitable for drug companies. Some scientists have suggested, with reason, that the success of these drugs has taken some of the impetus out of developing others. At the same time, the best way to make money in the drug business is by creating better drugs. Much of the effort in this area in recent years has gone into making serotonin-related drugs more specific, focusing on selective receptors rather than on presynaptic uptake in general. There are numerous serotonin receptors, and the more specific the drug can be made and still achieve a therapeutic effect, the fewer the side effects should be.

Because depression is not a single clinical condition, the same treatment doesn't work in all depressed people. One strategy being pursued is to find new treatments that involve novel combinations of drugs. For example, in some cases resistant to either tricyclics or SSRIs, a cocktail involving a mixture of an SSRI and a tricyclic is useful. The reason for this is not clear, but this is not unusual in psychiatric drug therapy, since so little is known about the underlying brain malfunction. In fact, one of the obstacles to the development of better drugs for psychiatric disorders has been that the main insights into the underlying brain malfunctions have come from findings, often accidental ones, about effective treatments. This is how the monoamine approach emerged, and why it continues to be a major focus of treatment strategies. However, there is growing recognition that, while monoamines are involved, the fundamental problem underlying depression is not in the monoamine system per se. Drug development is, as a result, moving in other directions, targeting other transmitters, modulators, hormones, second messenger systems, and neurotrophic factors. Some of these efforts are described below. As molecular biological approaches become better integrated with circuit-level approaches to brain function, new and better ideas about the nature of depression and its treatment are likely to emerge.

STRESS IS DEPRESSING

One of the most consistent biological findings about depression is that the adrenal cortex secretes more of the stress-related hormone cortisol in depressed people. This simple fact, which can be determined from a cotton swab containing saliva, links depression to the biology of stress. Over the years, a great deal of research has gone into understanding the relation between stress and depression, both through studies of depressed patients and experimental animal models of depression. However, recent findings about the adverse consequences of stress on the brain have led to a powerful new hypothesis about the origin of depression.

As we saw in chapter 8, in stressful situations, the concentration of cortisol rises in the bloodstream. This occurs because during stress, the amygdala and other brain regions alert hypothalamic neurons to release a peptide from their terminals in the pituitary gland. The peptide, called corticotrophin-releasing factor (CRF), causes the pituitary gland to release the hormone ACTH (adrenocorticotrophin hormone), which travels in the bloodstream and leads to the release of cortisol from the adrenal gland. The cortisol is then transported through the blood to various

bodily organs and tissues (it eventually reaches the salivary glands, which is why it can be measured in saliva). In the brain, cortisol binds to receptors in the hippocampus, among other regions. 'When a sufficient number of receptors are occupied in the hippocampus, signals are sent to the hypothalamus telling it to stop releasing GRE. In this way, the hippocampus regulates the stress response triggered by the amygdala, keeping the release of cortisol within a normal safe range.

In the short run, stress responses are useful in mobilizing bodily resources to cope with danger. But if the stress is severe and continuous, the consequences can be serious. Your cardiovascular system can be compromised, your muscles can weaken, and you can develop ulcers and become more susceptible to certain kinds of infections. But none of these should happen if the hippocampus is working properly to shut down the stress reaction. As we discussed in chapter 8, during prolonged and severe stress, the ability of the hippocampus to do its stress-control job falters.

Studies by Robert Sapolsky, Bruce McEwen, and others have shown that stress damages the hippocampus, leading to a shrinkage of dendrites and ultimately to cell death. Not surprisingly, functions that depend on the hippocampus, like explicit or declarative memory, become severely compromised. Stress hormones do not damage the hippocampus directly, but instead deplete its neurons of glucose, their main source of energy, and make them less capable of performing their job in the face of stringent demands. As a result, they become especially sensitive to elevations of excitatory transmitters like glutamate during periods of increased neural activity, such as occurs in stress.

Specifically, hippocampal cells have a toxic reaction to synaptically released glutamate in the presence of depleted glucose.

Cell shrinkage and cell death mainly take place in one area of the hippocampus the CA3 region. In another area, the dentate gyrus, cell death is far less likely to occur, but stress also takes its toll there. This is one of the few regions of the brain that is known to undergo neurogenesis, the production of new neurons, in adult organisms. Neurogenesis increases when animals learn new things, and decreases when they are under stress, or when they are given steroid hormones that mimic the elevations of cortisol that occur in stress. Together the cell shrinkage and cell death in the CA₃ area and the loss of new neuron production in the dentate gyrus probably account for the fact that the volume of the human hippocampus is smaller in people who have elevated levels of cortisol due to stress or other conditions.

For example, as we've seen, cortisol is elevated in depressed patients, who often have a smaller hippocampus and who accordingly suffer memory problems. Cortisol is also elevated in elderly people, especially those with memory problems and depression. Further, there's a condition known as Cushing's disease in which the adrenal cortex secretes excess cortisol. Cushing's patients also have a smaller hippocampus and memory problems, and many develop depression. Reduction of cortisol levels with drugs often reverses these structural and functional consequences. Similarly, people who receive prolonged steroid treatment for inflammation develop memory problems and depression, conditions that are reversed when the treatment is discontinued. One of the current trends in the treatment of depression is the development of drugs that alter the ability of GRE to lead to a rise in cortisol. Since cortisol is believed to be a major chemical culprit in stress (though perhaps not the only one), preventing cortisol from rising can sometimes help decrease stress and reverse depression.

The body system involved in controlling stress-hormone release is called the hypothalamic-pituitary-adrenal axis (HPA axis). This system did not evolve its ability to secrete cortisol for the purpose of making us depressed, but came about to help organisms deal with short-term physical dangers— situations that disrupt the normal physiological balance of the body, such as food or water deprivation, injury, or encounters with a predator or an enemy of your own species. The job of cortisol is to mobilize body resources in the short run, after which physiological normality, or homeostasis, is restored. As stress guru Robert Sapolsky notes, animals in the wild do not have to deal with thirty-year-long escaping sprints from lions, but humans have thirty-year mortgages and other long-lasting problems. It is the relentlessness of human stress, the long-term disruption of homeostasis, that increases our risk of disease,

including mental disorders like depression.

Ron Duman, a Yale researcher, has proposed a theory that explains how stress can predispose one to develop depression. According to Duman, depression results from an inability to make the appropriate adaptive response to stress. Duman's theory borrows heavily from the basic science findings of Sapolsky, McEwen, and others and emphasizes the effects of stress on the initiation of cell death in the CA₃ area and on the suppression of neurogenesis in the dentate area. In the most general sense, he proposes that stress elevates adrenal steroids and decreases neurotrophic factors, like BDNF, a molecule that is essential for cell growth and survival (chap. 4) and for the maintenance of synaptic connectivity (chap. 6). In the presence of increased neural activity and glutamate release, CA₃ and dentate cells depleted of glucose and lacking BDNF are endangered. In CA₃, they atrophy and die. In the dentate, new cell growth and connectivity of new cells with existing networks is suppressed. Following antidepressant treatment, though, these changes can be overcome—dendrites of CA₃ cells expand and the cells survive, and dentate cells continue to be generated and inserted into networks.

How, then, does antidepressant treatment achieve these goals? It has long been known that antidepressant drugs do not treat depression by simply altering synaptic levels of monoamines, since, as we saw earlier, their effects on synaptic levels are rapid but the treatment effects themselves take weeks. Duman proposes that the key to the drug effects are to be found in the second messenger cascades activated by the drugs. For example, when an SSRJ increases the amount of serotonin available at serotonin receptors, these receptors are stimulated for longer periods of time. As a result, a stronger intracellular response is generated and a more intense activation of second messenger systems occurs, leading to enhanced gene activation and protein synthesis. The long-lasting effects of drugs, then, may be due to new growth processes, which take time to be put in place and to alter synaptic networks.

Let's look a little more closely at how this might work. Recall from earlier chapters that during memory formation, calcium elevation inside cells raises cAMP levels, which activates protein kinases (PKA and MAP kinase) that phosphorylate the transcription factor CREB. CREB then induces genes to make proteins that become new tools of synaptic transmission (new receptors, new channels that pass calcium or other ions, new transcription factors or kinases, and so on). Antidepressant treatment skirts real-world learning experience and directly increases the amount of calcium inside the cell, triggering other second messengers that induce genes to make proteins. A brain on antidepressants can be brought back from its state of isolation from the outside world and encouraged, even forced, to learn. The brain, in other words, is duped into being plastic by these treatments.

Antidepressant medication thus enhances plasticity and thereby makes the brain more adaptive, better able to overcome the endangered state into which it was placed by excess cortisol. While the treatment itself does not directly substitute for direct experience (it doesn't give you new memories), by placing the brain in a state in which new memory formation is facilitated, depressed persons may be able to learn new mental states and behaviors that override the modes that they had been locked into by depression. A therapist who understands the patient and the underlying effects of the drugs may be able to serve as a guide to the recovery process, helping the patient turn to new, positive life experiences, or to turn back to old ones, at just the right time of antidepressant treatment.

Before leaving the topic of stress and depression, it's important to point out that the hippocampus is not the only region of the brain that is altered in depressed people. Imaging studies in humans show changes in size and in functional activity in the prefrontal regions as well. Further, animal studies show that this brain region has abundant receptors for adrenal steroid hormones and, like the hippocampus, is implicated in the regulation of HPA axis function. The elevated cortisol that occurs in depression therefore also attacks the prefrontal cortex, which might help account for other cognitive changes associated with depression, including poor short-term (working) memory, distractibility, and altered decision-making and executive functions. Changes also occur in the amygdala. More studies are needed to determine the extent to which antidepressant treatment reverses structural and functional changes associated with depression in these areas.

Depression, like schizophrenia, is thus no longer simply viewed as a monoamine imbalance. It is instead believed to involve altered circuits that lock one into a state of neural and psychological withdrawal in which the brain's

ability to attend to, engage, and learn about the world is reduced. Any treatment that can reengage a person with the world is likely to help. Whether psychotherapy alone can achieve this will likely depend on how far the maladaptive changes have progressed, as well as on factors peculiar to the individual, such as genetic makeup and past learning experiences.

MOTHER'S LITTLE HELPER

Anxiety disorders are the most prevalent form of mental illness. Schizophrenia afflicts about 1 percent of the population and depression about 10 percent, but roughly 25 percent of all adults are said to suffer from some manifestation of clinical anxiety. Everyone is a little anxious at times, but an anxiety disorder is said to exist when day-to-day psychological functioning and interpersonal relationships are disrupted by symptoms such as worry, tension, sleep disturbances, irritability, and somatic complaints. Anxiety is not a single disorder but a spectrum of conditions that includes nonspecific or generalized anxiety, phobias, panic attacks, posttraumatic stress disorder, and obsessive-compulsive disorder.

Psychotherapy is more effective in treating anxiety disorders than any other group of mental problems. This is true both of generalized anxiety and of the more specific conditions. The most effective forms of psychotherapy for anxiety involve behavioral and cognitive-behavioral approaches, which help the patient cope with and eliminate the unpleasant symptoms that accompany and perpetuate anxiety.

Medicinal approaches to anxiety have also been valuable. Alcohol is the oldest and most widely used anxiety-reducing drug. Many unwind from the tensions of the working world by stopping off for "happy hour" at their favorite watering hole, or with a "cocktail hour" at home. But because alcohol is intoxicating and addictive, it is not a practical solution for severe and chronic anxiety.

The first widely used medical treatment for anxiety involved barbiturates, long-acting drugs that depress brain activity. At low doses, they reduce tension and anxiety without producing intoxication, and at higher doses, they induce sleep. However, these drugs are highly addictive, and the difference between a dose that induces sleep and one that is fatal (by suppressing breathing) is quite small. The development of a new class of anxiety drugs in the 1950s was thus welcomed.

Meprobamate was identified in the process of screening variations of mephenesin, a drug that relieved anxiety but had a very short-lasting effect. One method that drug companies used to test the value of antianxiety drugs was to give them to monkeys, which tend to be wild and aggressive in captivity. Meprobamate made them calm, less vicious. When given to humans, it reduced anxiety without producing drowsiness. This compound was then made available by prescription under the brand names Equanil and Miltown,

and came to be known, in the popular press, as "happy pills." Though initially believed to be a safe, nonaddicting replacement for barbiturates, later reports were less favorable. Addiction did result, and sometimes medical complications occurred when the drug was stopped.

The next generation of drugs for anxiety emerged when a chemist who had worked developing dyes in Poland moved to the United States and joined a drug company. He tested a number of the compounds he had worked with in his dye research in Poland and found that one of them, and only one, was effective in calming wild monkeys. Upon further investigation, it turned out that this compound was different from all the rest. Chlordiazepoxide, as the compound was called, became the first of a new class of drugs, the benzodiazepines.

The effects of chlordiazepoxide were further tested in rats using the conflict test. The rats were made hungry, and then were allowed to obtain food by pressing a bar. However, if they pressed while a light was on, they were shocked. The light itself came to be a fear-arousing stimulus and led to outward signs of nervousness (freezing, defecation, etc.). When treated with chlordiazepoxide, though, they pressed more and were less visibly afraid even while the light was on. The drug allowed them to earn their daily bread and to be less fearful in doing so—perfect effects for a treatment of human anxiety. The conflict test in rats has, in fact, become one of the key methods for

determining whether new compounds are likely to be effective in relieving anxiety in humans.

Chlordiazepoxide was sold as Librium. Then came diazepam, better known as Valium, and alprazolam, marketed as Xanax. By 1975, 15 percent of the U.S. population was said to have taken at some point a version of “mother’s little helper.” Though highly effective in relieving anxiety, and far less problematic than barbiturates and meprobamate, benzodiazepines are likewise addictive, and can be dangerous when mixed with alcohol. Judy Garland, for example, is believed to have died from this combination.

In contrast to most other drugs used to treat mental illness, all the antianxiety drugs work immediately. That is why a drink can calm your nerves, and why Valium and other benzodiazepines are useful in acute anxiety. In *Starting Over*, a romantic comedy from the late 1970s, the character played by Burt Reynolds is having a severe anxiety attack in Bloomingdale’s while a crowd gathers around him. Someone calls out for a Valium, and essentially the entire crowd reaches into their purses or pockets and pulls out a vial. Valium jokes are, in fact, still popular in comedies. Hugh Grant’s character in Woody Allen’s recent film *Small Time Crooks* kept running to the bathroom for a Valium when things didn’t go his way.

While benzodiazepines are also useful as sleeping pills on an as-needed basis, these drugs can produce temporary memory loss after waking. There are infamous stories of businessmen taking a benzodiazepine to sleep on an overnight trip to Europe and then waking up and not being able to remember why they had flown overseas. A new class of drugs, imidazopyridines, are more useful for inducing sleep. Ambien is a prime example. Though their chemical structure is somewhat different from benzodiazepines, they work similarly (though more specifically) at the chemical level. Because these drugs have a shorter life span in the blood, and because they are not muscle relaxants, they have fewer lingering effects the next day on cognitive and behavioral functioning.

The classic antianxiety drugs work primarily by facilitating GABA inhibitory transmission in the brain, making it harder for glutamate to elicit excitation at its postsynaptic receptors. This is true for alcohol, barbiturates, and benzodiazepines (and imidazopyridines). However, each achieves its effects differently. Recall from chapter 3 that when a neurotransmitter binds to its receptor, the receptor opens, allowing electrically charged chemical ions to flow from the extracellular space into the cell. In the case of GABA receptors, the flow of chloride ions makes the inside of the cell more positive, which means that more negatively charged ions have to flow into the cell through glutamate receptors to initiate an action potential. Barbiturates enhance inhibition by acting directly on GABA receptors, keeping the chloride channel open longer, allowing more positive ions to enter the cell. Alcohol has a similar effect, but at a different GABA receptor from the one affected by barbiturates. Benzodiazepines, however, work differently, as they have their own receptors, which are linked to GABA receptors. So when benzodiazepine receptors are occupied, linked GABA receptors bind GABA more easily. As a result, the same amount of GABA released from a GABA terminal will have a greater inhibitory effect on the postsynaptic cell. The benzodiazepines therefore work only at those sites where GABA is being released naturally. Imidazopyridines work like benzodiazepines, but at a more selective part of the GABA receptor, which is probably why they have fewer side effects. (Parenthetically, the fact that benzodiazepine receptors exist in the brain means that the brain also has a natural supply of benzodiazepines. It may be that calm, worry-free people have more of this chemical, and anxious people less.)

Because GABA and benzodiazepine receptors are dispersed throughout the brain, drugs that bind to them do not affect anxiety exclusively. Anxiety may well be generated by specific networks, but drugs that treat it affect the entire brain. This accounts for some of the side effects of antianxiety drugs, such as drowsiness and muscle relaxation. Drowsiness, of course, is not considered a side effect when the drug is used for sleep induction rather than for anxiety reduction while awake.

Newer treatments for generalized anxiety are also available. Though these drugs produce fewer side effects, they have the disadvantage of not working immediately. Most of them enhance serotonin function. BuSpar (buspirone), for example, stimulates serotonin receptors (specifically serotonin 1A receptors). SSRIs, which as we’ve seen make more serotonin available at receptors, also appear to be helpful with these disorders.

WORRY CIRCUITS

'Where does anxiety come from? Jeffrey Gray has long believed that the antianxiety drugs are the key to understanding the nature of generalized anxiety; and specifically the circuits in the brain that generate anxiety; As he argues, alcohol, barbiturates, and benzodiazepines are different drugs from a chemical point of view, and each has different side effects, but each relieves anxiety; By studying the effects of the three drugs on a variety of different behavioral tasks in rats, and identifying a set of tasks that are affected by all three drugs, a behavioral portrait of anxiety can be drawn. And if brain regions can be discovered that when damaged lead to the same effects as the antianxiety drugs on these "anxiety" tasks, then the neural home of anxiety might be isolated.

This research program led to a neural theory of anxiety published in 1982. Because the common effects of the antianxiety drugs were similar to the effects of damage to the septum and/or hippocampus, these areas were postulated to be at the core of anxiety. We've encountered the hippocampus frequently, but not the septum. This is a region of the forebrain—actually part of the so-called limbic system—that is intimately associated with the hippocampus and regulates some of its activity.

From a psychological point of view, the septum and hippocampus were thought of as constituting the brain's behavioral inhibition system, a network that detects and responds to aversive stimuli, those that produce pain, punishment, failure, or loss of reward, or that elicit novelty and uncertainty. 'When the behavioral inhibition system is activated, ongoing behavior is inhibited (resulting, for example, in freezing), and the organism becomes aroused, attentive, and vigilant. Administering antianxiety drugs prevents aversive stimuli from eliciting septohippocampal arousal and vigilance, and damaging these areas reduces anxiety by eliminating the brain regions where arousal and vigilance would be elicited by aversive stimuli.

As we saw above, the classic antianxiety drugs work by enhancing GABA transmission. Although the drugs each affect GABA transmission by different means, they all work through GABA. A key part of Gray's theory was that GABA neurons were the common denominator in anxiety. However, because classic antianxiety drugs, when administered to the hippocampus itself, proved not very effective in relieving anxiety, Gray proposed that these substances alter GABA transmission in another brain region that subsequently influences the septohippocampal system. The key area of action was proposed to be the monoamine systems in the brain stem, specifically the serotonin and

norepinephrine systems. These areas were selected on the basis of earlier studies showing that, when stimulated in rats or monkeys, they elicited anxiety-like behaviors, and that antianxiety drugs injected in the same areas inhibited cell activity and relieved anxiety—like behaviors.

Before considering Gray's theory of how antianxiety drugs work, let's briefly look more closely at his conception of what anxiety is at the level of neurons. During threats, serotonin and norepinephrine cells in the brain stem are activated, and serotonin and norepinephrine are released from the terminals of these cells. Although the terminals are located in many regions of the brain, because serotonin and norepinephrine are modulators, their main effect is to alter transmission at active synapses rather than to elicit activity themselves (chap. 3). Since the septum and hippocampus are involved in processing threats, these regions would be active in the presence of threats, and serotonin and norepinephrine would enhance their synaptic processing, leading to arousal and vigilance, and anxiety.

Now for the antianxiety effects. The antianxiety drugs, as we've seen, enhance GABA transmission. The enhancement mainly occurs at synapses where GABA transmission is taking place, rather than at inactive synapses. Because GABA activity is often triggered in response to excitation in nearby cells, the enhancement of GABA by antianxiety drugs in the presence of threatening stimuli will be concentrated at those synapses processing the threat. In addition to activating cells in the septum and hippocampus, threats activate serotonin and norepinephrine cells. Gray therefore proposed that antianxiety drugs enhance GABA transmission at these brain stem cells, and thereby reduce the release of serotonin and norepinephrine in the forebrain. The enhanced processing that these chemicals were engendering in the septum and hippocampus will in turn be eliminated, decreasing arousal and vigilance, thereby reducing anxiety.

Although Gray's theory was a tour de force of neuropsychological logic, it ran into trouble. First of all, it

focused exclusively on GABA transmission, since this was a point of convergence between all the classic antianxiety drugs. The emergence of drugs that successfully treated anxiety without affecting GABA transmission therefore challenged the theory. Buspirone and SSRIs, for example, are both effective in relieving anxiety and both work at the level of receptors postsynaptic to serotonin terminals rather than at cell bodies in the brain stem. Indeed, these drugs are somewhat effective when injected directly into the hippocampus, where they selectively enhance serotonin transmission. While this helps salvage part of the theory, it weakens the core concept that facilitation of GABA transmission in the brain stem is key. Second, the theory did not allow a role for the amygdala, which as we've seen plays an important role in processing danger and threats.

In a recent revision of the theory, Gray and his longtime collaborator Neil McNaughton attempted to address both of these shortcomings. Their new formulation conceives of anxiety as a psychological state emerging out of a synaptically connected network involving, among other regions, the septohippocampal area, amygdala, and prefrontal cortex, all modulated by brain stem monoamine systems. Although the new theory is broader in its coverage, in my opinion, it still gives the septum and hippocampus too prominent a role, at the expense of the amygdala and prefrontal cortex.

Anxiety, in my view, is a cognitive state in which working memory is monopolized by fretful, worrying thoughts. The difference between an ordinary state of mind (of working memory) and an anxious one is that, in the latter case, systems involved in emotional processing, such as the amygdala, have detected a threatening situation, and are influencing what working memory attends to and processes. This in turn will affect the manner in which executive functions select information from other cortical networks and from memory systems and make decisions about the course of action to take.

I believe that the hippocampus is involved in anxiety not because it processes threats, as Gray suggests, but instead because it supplies working memory with information about stimulus relations in the current environmental context, and about past relations stored in explicit memory. When the organism, through working memory, conceives that it is facing a threatening situation and is uncertain about what is going to happen or about the best course of action to take, anxiety occurs.

In my formulation, the amygdala plays a greater role in threat processing than the hippocampus. I have described the role of the amygdala in processing and responding to threats in earlier chapters and in *The Emotional Brain*, and will only briefly recount it here. When sensory information about a threatening stimulus is detected by the amygdala, output connections to response-control systems in the brain stem initiate the expression of defense responses (freezing) and supporting physiological changes in the body (rises in blood pressure and heart rate, stress hormone release, and so on), some of which give rise to signals that are fed back to the brain and influence ongoing processing. Although Gray's theory did not address the question of how monoamine systems are activated by threat, one of the key ways this occurs is through direct connections from the amygdala to the monoamine cells. Thus, during threats, serotonin and norepinephrine (and dopamine) are released in widespread forebrain areas (including the prefrontal cortex, hippocampus, and amygdala, as well as other areas). Direct connections from the amygdala to prefrontal areas (including anterior cingulate and orbital cortex) allow the detection of threats by the amygdala to directly influence working memory processing. But the prefrontal cortex and its working memory functions are also influenced through other routes, including the amygdala-initiated release of monoamines and the feedback from hormonal and other bodily responses. When the amygdala detects a threat, it triggers consequences that ultimately place working memory in a vigilant processing state, causing it to continue to attend to whatever it is occupied with at the moment, biasing thoughts, decisions, and actions. Further, output connections of the amygdala to the nucleus accumbens, as we saw in the last chapter, allow threatening stimuli to motivate the organism to avoid the source of the threat. This is an important factor, given that pathological avoidance of possible threat sources is a paramount behavioral symptom of anxiety disorders.

Up to this point in the discussion, I've not made a clear distinction between fear and anxiety. Classically, though, fear is viewed as a reaction to a specific and immediately present stimulus, whereas anxiety is a concern

about what *might* happen. One possible resolution of the Gray-McNaughton theory with the amygdala theory would be if amygdala networks took care of fear and hippocampal networks took care of anxiety. This might have been an acceptable thesis some years ago, but there is now strong evidence showing that injection of benzodiazepines directly into the lateral and basal amygdala, the input stages of the amygdala, reduces anxiety behaviors in several (though not all) of the classic tasks used to test the efficacy of antianxiety drugs. That finding is consistent with the fact that benzodiazepine receptors are concentrated in these input regions of the amygdala. Further, atypical antianxiety drugs that target serotonin receptors also relieve anxiety in animal models when injected into the amygdala. Both classical and atypical antianxiety drugs may therefore achieve their anxiety-reducing effects, at least in part, at the input stages of amygdala processing, making it harder for threatening stimuli to initiate activity in the amygdala, thereby preventing amygdala activation from arousing the rest of the brain.

But another theory has been proposed that distinguishes fear and anxiety in a different way. Michael Davis has performed experiments suggesting that anxiety might be a function of the bed nucleus of the stria terminalis, a brain region that is considered an extension of the amygdala and whose output connections are remarkably similar to those of the amygdala. Because of their similar outputs, the amygdala and bed nucleus can affect many of the same target brain areas (prefrontal cortex, monoamine systems, and so on) in comparable ways and produce the same kinds of bodily responses (muscle tension, fast-beating heart, sweaty palms, tight stomach, and so on). But because the inputs to the two structures are different, they might be activated under different conditions—the amygdala in response to immediately present threats, the bed nucleus to anticipated ones. This distinction between the role of the amygdala in fear and the bed nucleus in anxiety is appealing because it potentially explains why people on antianxiety medication can be generally less worried (that is, less anxious), and still be capable of responding to an immediate threat. However, much work remains to be done to clarify the role of these two regions and their relation to the broader circuits involved in anxiety.

In summary, generalized anxiety is an aroused state of mind initiated and maintained by emotional processing. As a result, it requires, at a minimum, networks involved in arousal (monoamine systems), emotional (amygdala, perhaps including the extended amygdala), and cognitive (prefrontal cortex, hippocampus) functions. And while individual brain regions and networks make distinct contributions to the processes that together constitute anxiety, anxiety itself is best thought of as a property of the overall circuitry rather than of specific brain regions.

THE SPECTRUM OF ANXIETY

So far, this overview has focused on free-floating or generalized anxiety. But the anxiety disorders also include phobias, panic disorder, posttraumatic stress disorders, and obsessive-compulsive disorder. These are distinct conditions with different causes and symptoms, and the treatments that work for one do not necessarily work for the others.

Benzodiazepines are of little help in treating the so-called simple phobias (pathologic fear of snakes, spiders, heights, and so forth). In fact, the most effective treatment for these conditions is psychotherapy, especially cognitive behavioral therapy, in which the symptoms are reduced by exposure to the feared stimulus under various conditions (see the box entitled “Cognitive Behavioral Therapy”). Another kind of phobia, social phobia, centers around a fear of being negatively evaluated by others—for example, by people with whom one works or shares interests. Some forms of social phobia involve a fear of being the center of attention in public. Examples of the latter include an extreme debilitating fear of public speaking or performance, and fear of eating in public or using public bathrooms. Antidepressant medications, including certain SSRJs and monoamine oxidase inhibitors, are often used in treating these conditions. Cognitive behavioral therapy is also useful. Drugs like propranolol, which reduce the bodily signs of fear (high heart rate and so forth) but that do not affect the brain directly, are sometimes used to aid in public speaking and stage fright. As with social phobia, the preferred treatment for panic disorder and posttraumatic stress disorder today is the use of an SSRJ in combination with some form of psychotherapy involving anxiety management and cognitive behavioral therapy. A variety of other drugs (including monoamine oxidase inhibitors, tricyclics, and

benzodiazepines) also

have some positive effects in these disorders, but the SSPIs have a better side-effect profile. Obsessive-compulsive disorder likewise responds to SSRIs to some degree.

Because the different forms of anxiety are manifested in different mental, behavioral, and physiological symptoms and often under distinct conditions (having to give a speech may have little effect on a spider phobic but will be a source of fear and anxiety to one who has a fear of public speaking), there is something unsatisfactory about the use of similar medications to treat them all. And because the preferred medications are antidepressants, rather than drugs developed for specific anxiety disorders, one can't help but feel that from the point of view of drug development and treatment, we have hardly touched these mental problems some of which are widespread and devastating. Certainly, following the strategy of using information about effective treatments to give clues to the underlying biological disturbance is not going to reveal differences between various anxiety disorders, since the same treatment (SSRIs) is used for different anxiety disorders, as well as for depression.

Earlier, I mentioned that I am the director of the Center for the Neuroscience of Fear and Anxiety (<http://www.cns.flyU.CdU/CNFA/>). This center brings together scientists at all of the major biomedical research institutions in Manhattan, including New York University (Elizabeth Phelps, Mony de Leon, and me), Rockefeller University (Bruce McEwen), Cornell University Medical College (David Silbersweig and Emily Stern), Mt. Sinai Medical School (John Morrison and Patrick Hof), and Columbia University (Jack Gorman). The research involves studies of animals and of patients with anxiety disorders. Our immediate goal is to elucidate the ways in which the brain is changed in anxiety disorders, in the hope of facilitating, over the long run, the development of more selective and more effective treatments with fewer side effects. The patient work involves the use of state-of-the-art functional MM imaging techniques and the animal work takes advantage of the vast arsenal of neuroscience tools that has been discussed throughout this book.

Because anxiety disorders involve alterations in the processing and/or reaction to threat and danger, and fear mechanisms are to a large extent conserved in the brains of humans and other mammals, research on the psychology and neuroscience of fear can inform ideas about the underlying disturbance in different anxiety disorders. Although there have been several imaging studies of patients with anxiety disorders of various types, the aim of our center is to use the enormous progress that has been made in elucidating the neural pathways underlying fear conditioning in animals to try and understand the nature of the brain changes that take place in anxiety disorders, especially panic disorder and posttraumatic stress disorder (PTSD).

Panic disorder and PTSD both involve alterations in the processing of threats, but the causes and symptoms are distinct. In panic disorder, fear is expressed in the form of discrete and sudden panic attacks that arise without an obvious environmental stimulus and lead to chronic anticipatory anxiety and various degrees of avoidance behavior. Fear in panic patients has no apparent relationship to any actual threat and often involves abnormal sensitivity to uncomfortable somatic sensations. In PTSD, fear is expressed as increased sensitivity to stimuli reminiscent of an original, life-threatening traumatic event and includes flashbacks and increased startle response. Different neural mechanisms should therefore be altered in these patients.

As we've seen repeatedly, the amygdala is a key brain region in the acquisition and expression of fear elicited by stimuli associated with threatening experiences. And several studies of normal humans have now shown amygdala activation during fear conditioning, even when the conditioned stimulus is presented subliminally, that is, unconsciously." We therefore can use fear conditioning to activate the amygdala and determine whether the pattern of activation there is altered in patients with anxiety disorders. Our expectation is that because changes in the way fear-arousing stimuli are perceived and responded to characterize all anxiety disorders, we should see increased levels of activity in the amygdala in both panic and PTSD patients.

But the amygdala is not the only area implicated in fear conditioning. As we've seen, the hippocampus is involved in processing fear-arousing situations or contexts, and the medial prefrontal cortex in the adjustment of fear reactions in response to changing environmental conditions, as in extinction (the process whereby a conditioned

stimulus loses its fear-arousing properties when it no longer predicts danger). Recent studies by Liz Phelps have confirmed the role of the human hippocampus in processing contextual fear stimuli, and we are currently testing whether activation of the human medial prefrontal cortex will be altered during extinction.

The use of fear conditioning allows us to assess whether patterns of neural activity within and between the amygdala, hippocampus, and medial prefrontal cortex differ in panic and PTSD patients from normal controls. Further, we can determine whether successful treatment (e.g., the use of SSRIs and cognitive behavioral therapy) will cause the brains of these patients to look more like those of normal controls. By testing patients after treatment

with either SSRIs or cognitive behavioral therapy alone, we should be able to determine whether the two kinds of therapy achieve their results by altering the same or different components of the circuits.

Once we have some sense of how the brains of panic and PTSD patients differ from each other and from those of normal persons, we can turn to animal studies and explore what kinds of life experiences and biological conditions might predispose the brain to change in these ways. For example, suppose we find that the amygdala is hyperactive in both panic and PTSD, but that the hippocampus and prefrontal cortex change in different ways in each condition. We could look for stressful experiences that produce these patterns of change in the rat brain. If such patterns can be identified, then the changes can be pursued at the level of synaptic alterations and explained in biological terms: as neurotransmitter actions on receptors that open channels and allow calcium to flow inside the cell and trigger kinases that phosphorylate transcription factors that induce genes to make proteins that stabilize synaptic changes. To the extent that any one of these molecular steps differs between anxiety disorders, clues about future drug development would be available.

By now you've surely noticed a trend. Neuroscientists have proposed that the prefrontal cortex, hippocampus, and amygdala are altered in some way in all forms of mental illness we've considered so far. It seems that mental disorders are giving us a message, telling us that these three brain regions are especially critical to understanding who we are and why we are that way. But we should remember that the same brain areas can be involved in different disorders for different reasons. At a fairly gross level of explanation, neural activity in a particular area might be up in one disorder and down in another. But in order for neural activity to change in this way, more basic changes have to be taking place within and between individual cells in the area in question. Since these deeper changes can be revealed only by exacting neurobiological studies, which can be done only in experimental animals, it is essential that research on psychiatric disorders in humans and basic science investigations of brain function in animals be coordinated.

We have just begun this research program, and it is too early to know just what we'll turn up. However, because of the systematic nature of the research program, we will almost certainly be able to point to differences in the brains of panic and PTSD patients, even if they are not the ones we expected. But, in many ways, unexpected findings are the most interesting, as they force thinking about problems in new ways.

BRAINS GENES, AND THERAPY

We've gotten all the way to the end of this chapter on biological approaches to mental illness without discussing genes. To many, especially to critics, a biological approach to mental disorders implies an emphasis on genetics. But I hope I've presented a credible alternative view. As in all other aspects of mental life, genes can predispose people in certain directions without altogether predetermining their destiny. Genes contribute to, rather than solely dictate, synaptic connectivity.

Critics such as Cohn Ross and Alvin Pam are concerned that a biological approach to mental disorders implies that "the individual with behavior problems must suffer in some way from defective protoplasm, a constitutional predisposition to mental illness," and that the body or brain is responsible for mental illness rather than the family or society. They strongly equate biological psychiatry with genetics. However, while a genetic approach to mental disorders is by definition biological in orientation, the reverse is not necessarily true: a biological approach is not exclusively or even primarily genetic. Let's first take a look at the genetic approach, and then consider the broader

picture.

A massive effort is now under way to determine the extent to which genes contribute to mental illness. For example, years of research have gone into pedigree studies of families or groups of people with a tendency toward a mental disorder, and into concordance studies of identical twins, which determine the likelihood that the presence of the disorder in one predicts the disorder in the other. Schizophrenia, depression, anxiety disorders, alcoholism, and a variety of other psychiatric and neurologic conditions are being studied this way.

For illustrative purposes, consider the work on schizophrenia. The research shows a strong correlation between the number of genes shared and the likelihood of developing the disorder. Schizophrenia occurs in about 1 percent of the general population. In contrast, in identical twins (100 percent gene overlap), if one child has schizophrenia then there's roughly a 48 percent chance the other will develop it at some point. But in fraternal twins (50 percent gene overlap), the likelihood drops to 17 percent. In siblings (25 percent gene overlap), the figure falls to 9 percent, and in first cousins (12.5 percent gene overlap) to 2 percent. The overall picture that emerges is that schizophrenia is strongly tied to genetic factors. That is, in genetic terms, there's a concordance among identical twins. But one can view this cup as half full or half empty—there's also a discordance. If genes fully “explained” schizophrenia, concordance would be 100 percent.

The discordance between identical twins is due to several factors. One is that gene expression is an *epigenetic* phenomenon—it involves the interaction of the gene with environmental factors. An extreme example of epigenesis is the disease phenylketonuria (PKU), a genetic form of mental retardation that is not expressed unless the essential amino acid phenylalanine is ingested. (It is present in dairy products, avocado, certain nuts, and in the sugar substitute aspartame.) In this case, early detection and treatment can prevent the symptoms of the disease. In most genetic conditions, though, the environmental co-conspirators are not known. Another factor working against concordance is the polygenetic nature of the genetic contribution. In the days when psychiatric disorders were thought to be due to a simple imbalance in a particular brain chemical, it was plausible that the disorder might be caused by a single gene (one that makes transmitter or receptor molecules). But as we've seen, psychiatric disorders involve complex circuits in the brain, and therefore any genetic contribution is likely to involve interactions among multiple genes. Because the expression of every gene also involves interactions with environmental factors, polygenetic inheritance greatly expands the manner in which the environment can influence genetic predispositions. Further, in some cases, people with the gene for a disease show little or no evidence of having the symptoms of the disease. This is called *nonpenetrance*. In other instances, genetic diseases are variably expressed, where some family members may have the full-blown disease and others only mild symptoms. Finally, mental disorders can probably have multiple unrelated causes. It is believed, for example, that in some instances, schizophrenia can be induced by environmental factors, like brain injury or infection, in persons with little or no genetic leaning toward the disorder. So discordance between identical twins might sometimes be due to the fact that the afflicted twin acquired schizophrenia for nongenetic reasons.

Biological psychiatrists are interested in the possibility that genes contribute to mental disorders because genetic problems might in the future be fixable. A gene working overtime and thereby causing a malfunction might, with a pill, be turned down or off, and one that is sluggish might be encouraged to work harder. Whether these feats will ever be successfully accomplished is unknowable at the present. But if they are, many people are likely to be helped.

It should be noted that success or failure for one psychiatric disorder will not predict the success or failure in treating others. Each disorder will have its own peculiar advantages and disadvantages for a genetic approach, and each has to be attacked individually. At the same time, research strategies developed for one problem may help in the pursuit of the others. With the human genome recently pronounced as “knowable,” this kind of work is likely to expand in the coming years.

But perhaps I'm not being completely fair to the critics. One of their concerns is that a focus on genes or other aspects of biology puts too much blame on the body, discouraging mental professionals from using nondrug

forms of therapy, insurance companies from paying for nondrug therapy, and patients from taking responsibility for their own healing. But these concerns, I believe, reflect a misunderstanding of the aims of biological psychiatry. In spite of the current enthusiasm for genetic analyses of mental disorders, most biological psychiatrists view genes in a realistic way. Raymond DePaulo, a biological psychiatrist at Johns Hopkins University, expresses what is probably the most common view: "I don't think depression is all genetic. If you can nail down one end of the tent, however, it's easier to figure out by stretching out the rest of the tent what it should look like." In other words, once you know the genetic component of a disorder, it should be easier to figure out which kinds of environmental experiences interact with it to trigger the disorder and make it worse. Genes are part of the story, and if that part of the story can be interpreted, we should be better off than if we ignore it.

Experience is often considered the counterpoint to genes. But *experience* is a complex notion with many possible meanings and endless implications. As a result, the role of experience is probably going to be even more difficult to understand than the role of genes. But the more we understand about how experience changes the brain, the better off we'll be in discovering how it contributes to mental disorders. As we've seen, experiences affect the brain when they are stored as synaptic changes in one or more systems during learning, which is why research on synaptic plasticity, and on learning and memory, is so important.

Where does a biological approach to mental health and mental disorders leave psychotherapy? Eric Kandel, the renowned neuroscientist, started his career as a psychiatrist. Recently, he boldly proposed a "new intellectual framework for psychiatry," one founded on biological principles. He argued that we are on the threshold of understanding memory and emotion systems in the brain, and that this information, combined with continued advances in molecular neuroscience, offers a wealth of opportunities for psychiatry.

As I pointed out at the beginning of this chapter, psychotherapy is fundamentally a learning process for its patients, and as such is a way to rewire the brain. In this sense, psychotherapy ultimately uses biological mechanisms to treat mental illness. This does not mean, however, that psychotherapy involves learning while drug therapy involves something else, like correction of genetically dictated chemical imbalances. Even if chemical imbalances were to account for mental disorders, the imbalance could result purely from environmental factors, like some intensely stressful experience, or from environmental events that trigger and amplify a genetic predisposition. The therapist's job, whether using drugs or not, is to restore mental well-being. If the patient's problem involves neural changes that are locking the brain in an aroused state, relief will come only when the aroused state is reduced, either by psychotherapy or drug therapy, or by some combination of the two. Regardless of whether the initial cause is social stress or a genetic time bomb, unless the changes in the brain that accompany the disorder can be reversed or circumvented, the problem is unlikely to dissipate.

Contemporary approaches to the treatment of mental disorders are not distinguished from older ones so much by an emphasis on drugs as by an emphasis on *effective* treatment. Jack Gorman of Columbia University, a leading biological psychiatrist, puts it this way: "In the Old Psychiatry, each clinician learned a set of techniques from one school of thought and insisted on applying them to all corners. In the New Psychiatry clinicians may still feel most comfortable with the tools of one school, but they readily admit when they don't have the right skill for a particular patient. Today, psychoanalysts are willing to refer patients who need medication to a psychiatrist with skills in psychopharmacology, just as those trained in psychopharmacology increasingly send patients to a psychotherapeutic specialist. "The New Psychiatry," according to Gorman, "is only interested in providing safe and effective treatment. If that means borrowing from different points of view to offer a patient a combination of therapies, then that is what must be done."

Once a disorder exists, and the brain has changed, the changes have to be dealt with in some way in order for a patient to recover. Drugs can induce adaptive changes in neural circuits, or put neural circuits in a state where adaptation and learning are promoted. But there's no guarantee that, left to its own devices, the brain will learn the right things. Patients, in other words, are likely to benefit most from drug therapy when the drug-induced adaptivity of their brains is directed in a meaningful way. This is probably best achieved by traveling down the pharmacological

road to recovery with some one who understands not just the drug or the person, but the drug, the person, and the life situation the person is experiencing. HMOs may not like it, but the drug, therapist, and patient are partners in the synaptic adjustment process called therapy, with drugs attacking the problem from the bottom up, the therapist from the outside in, and with the patient reaching up and down from his or her own synaptic self.

CHAPTER ELEVEN

WHO ARE YOU?

WHO ARE YOL—WHO WHO, WHO WHO?

—The Who

Figuring out how the brain works is a daunting task. For that reason, neuroscientists usually work only on pieces of the puzzle—like aspects of cognition, emotion, or motivation—rather than on the whole organ and its systems at once. But if we want to understand how our brains make us who we are, we have to figure out how these individual processes blend together to cause a person to emerge effortlessly from the electrochemical activities of the protoplasmic mass that is his or her brain. It's time for me to bite the bullet and explain how I think the brain, specifically its synapses, makes us who we are.

BRAINS AND OTHER PARALLEL COMPUTERS

I once had a friend who supported a bohemian life in downtown Manhattan with a job as a computer programmer in a neuroscience lab. He was extremely talented, too talented for such work, and eventually got a position more suited to his abilities. He was hired by a company in Boston that made powerful computers called "connection machines." It was through him that I first heard about parallel computers.

Parallel computers work differently from the standard model with which we're familiar. Rather than doing computations one at a time in sequence (in other words, carrying out the steps of a program line by line in serial order), they process many steps simultaneously. Parallel computers can function this way because, in contrast to your desktop PC or Mac, they have many processing units that can be devoted to the execution of a given task. By distributing the workload across the various processors, they can perform the task much faster than serial computers. (Don't plan on getting one for your home or office, however: they cost millions of dollars.)

The brain is also sometimes described as a parallel computer, but it actually functions differently from an off-the-shelf connection machine. The brain is organized into processors (neural systems) that function independently of one another (at least to some extent). Since each of these systems has a specific job assignment, several types of tasks can be done by the brain simultaneously, that is, in parallel. This architecture enables you to chew gum and walk down the street, guiding yourself toward your destination while feeling happy and rehearsing the phone number

your friend gave you a block back, all at the same time as your posture is maintained upright, your blood pressure is kept at a safe level, and your rate of breathing is paced to the oxygen needs imposed by all the activities in which you're engaged.

Connection machines can, like brains, be divided up in such a way that different groups of processors are responsible for particular tasks. Although each task is then performed less efficiently than it would be if all the processors were devoted to it, overall this can be a more efficient use of the machine since multiple tasks can be worked on at the same time. Reversing the logic, if we had fewer neural systems using up the same overall computing power in our brains the systems would each be more powerful. However, because we have to do lots of different things each day to stay alive and well (eat, sleep, walk, avoid danger and pain, hear, see, smell, taste, talk, and think, to name some), with fewer brain systems we would almost certainly be less capable, even if the remaining systems were each more proficient at their particular tasks.

Over the long, slow evolutionary process of building the brains of vertebrates, and then of mammals, and eventually of our own species, the neural systems we have were specifically designed to take care of important jobs. We don't have extra ones that we can easily give up and continue to lead life as before. Similarly, new systems aren't easily acquired. For example, the addition of language and related aspects of cognition to the primate brain was not a trivial process—the brain was by that point in its evolution already fully booked. Adding a whole new set of functions, therefore, required either a loss

of space devoted to some functions or an increase in brain size. In fact, both seem to have occurred. The human brain is bigger than that of other animals (relative to our body size) and also seems to have undergone some reorganization. For example, the neural mechanism underlying the perception of spatial relations is present in both hemispheres of other primates; it is mainly on the right side in humans. This implies that spatial perception was forced from the left during the language invasion of human synaptic territory.

Life requires many brain functions, functions require systems, and systems are made of synaptically connected neurons. We all have the same brain systems, and the number of neurons in each brain system is more or less the same in each of us as well. However, the particular way those neurons are connected is distinct, and that uniqueness, in short, is what makes us who we are.

THE PARADOX OF PARALLEL PLASTICITY

We've met a number of the brain's neural systems throughout this book. Included are networks involved in sensory function, motor control, emotion, motivation, arousal, visceral regulation, and thinking, reasoning, and decisionmaking. What is remarkable is that synapses in all of these systems are capable of being modified by experience. Consider a few examples.

Emotion systems, as we've seen, are programmed by evolution to respond to some stimuli, so-called innate or unconditioned stimuli, like predators or pain. However, many of the things that elicit emotions in us or motivate us to act in certain ways are not preprogrammed into our brains as part of our species heritage but have to be learned by each of us. Emotion systems learn by association—when an emotionally arousing stimulus is present, other stimuli that are also present acquire emotion-arousing qualities (classical conditioning), and actions that bring you in contact with emotionally desirable stimuli or protect you from harmful or unpleasant ones are learned (instrumental conditioning). As in all other types of learning, emotional associations are formed by synaptic changes in the brain system involved in processing the stimuli. Some of the brain's plastic emotional processors include systems involved in detecting and responding to danger, finding and consuming food, identifying potential mates and having sex.

Sensory systems are likewise plastic. Until very recently, it was believed that perceptions are stable from day to day and year to year because sensory systems are immutable after childhood. But studies by Norman Weinberger, Charles Gilbert, Mike Merzenich, and others have shown that these systems are remarkably susceptible to modification in response to external stimulation. Merzenich, for example, has shown in a variety of ways how experience with a particular stimulus alters the net area, and thus the underlying synaptic territory; devoted to processing that stimulus. Following amputation of a finger, for example, the area in the somatosensory cortex devoted to that finger shrinks, while extensive stimulation of a particular finger expands the area of cortex devoted to it.

Motor systems are also plastic. This is obvious from the fact that we can learn skills and can improve our ability to perform certain movements with practice. As we saw in chapter 5, changes in synaptic transmission in the cerebellum are important for some forms of motor skill learning. Synaptic plasticity also occurs in the motor cortex, basal ganglia, and other brain regions involved in motor control.

Because synaptic plasticity occurs in most if not all brain systems, one might be tempted to conclude that the majority of brain systems are memory systems. But as I argued in chapter 5, a better way of thinking about this is that the ability to be modified by experience is a characteristic of many brain systems, regardless of their specific function. Brain systems, in other words, were for the most part not designed as storage devices—plasticity is not their main job assignment. They were instead designed to perform particular tasks, like processing sounds or sights, detecting food or danger or mates, controlling actions, and so on. Plasticity is simply a feature that helps them do their job better.

The fact that plasticity does occur in so many brain systems, however, raises interesting questions. How does a person with a coherent personality—a fairly stable set of thoughts, emotions, and motivations—ever emerge? Why don't the systems learn different things and pull our thoughts, emotions, and motivations in different directions? What makes them work together, rather than as an unruly mob?

LESSONS FROM DISCONNECTION

Before we examine what holds the self together, let's consider how fragile a patch job it is. The bottom line is simple: Functions depend on connections; break the connections, and you lose the functions. This is true of the function of a single system (a lesion of the visual thalamus, for example, will prevent information from the eyes from reaching the cortex, and thus will prevent the cortex from being able to perceive the visual world) as well as of interactions between systems (a lesion in a certain part of the temporal lobe will prevent information about visual objects from reaching the prefrontal cortex and thus will prevent that stimulus from being held in working memory, and hence from being used as the basis for thinking and decision-making).

The most striking instance of a broken connection I've ever personally witnessed involved a teenage boy who underwent split-brain surgery for the control of epilepsy. In this operation, the nerves between the two sides of the brain are severed to prevent the seizures from moving between the left and right hemispheres. One of its unwanted consequences is that the two hemispheres become somewhat independent. Several days after the surgery, the boy was observed pulling his pants down with his right hand and up with his left. Since the right hand is under the control of the left hemisphere, and the left hand of the right hemisphere, the normal integration of motor control that exists so effortlessly in each of us broke down in this boy because the brain systems that had final say over what the two hands do were no longer connected. (This sort of conflicted behavior has been described for many such patients but was amazing to see firsthand.)

Another example involves conduction aphasia. Patients who suffer from this condition can speak without trouble and can understand spoken words (for example, they can point to the picture of an object named by someone else), but cannot, upon hearing a spoken word or sentence, repeat it, or answer a question posed to them. The reason for this, according to the late Norman Geschwind, one of the great neurologists of the twentieth century, is that the neural pathway that transmits information between the areas of the brain responsible for the comprehension and production of speech is cut. Conduction between the brain areas is thus disrupted.

Geschwind used the phrase *disconnection syndromes* to characterize conditions in which specific behavioral or mental consequences result from disrupting communication between brain regions. What is striking about these disorders is that the deficit results not from the loss of a particular function, but from the inability to exchange information between brain areas. Disconnection syndromes illustrate how critical internal coordination between brain systems is in maintaining the unity of mind and behavior.

While the effects of brain lesions are not invariably interpreted as disconnection syndromes, disconnections always occur when the brain is damaged. For example, lesions of the prefrontal cortex, especially the ventral or orbital frontal cortex, have, since the nineteenth century been known to drastically alter personality. After an iron rod passed

through the head of Phineas Gage in a railroad accident, this upstanding citizen suddenly became fitful, irreverent, and unrestrained. As one contemporary observer put it, "The equilibrium between his intellectual faculties and animal urges seems to have been destroyed." Building on this case and numerous others, Antonio Damasio has suggested that damage to the ventral prefrontal cortex can lead to a loss of social control, and in the extreme can cause sociopathic behavior. But it would be wrong to think of the ventral prefrontal region as the center of social grace. According to Damasio, the consequences of damage to the ventral prefrontal cortex can be conceived of as a breakdown in the ability to use emotional information to guide thoughts and actions. This seems reasonable given that some of the key connections of the ventral prefrontal cortex include areas involved in higher cognitive processes (other prefrontal regions such as the lateral prefrontal cortex and anterior cingulate cortex, as well as the hippocampus) and emotional and motivational functions (the amygdala and nucleus accumbens). Ventral prefrontal damage thus does more than just create a hole in this part of the brain. It removes that brain region from the circuits in which it participated. Regardless of whether it is obvious or not, brain lesions always produce disconnections.

Connectivity changes also take place in psychiatric disorders, but they are typically more subtle than those in neurological patients with overt brain le

sions. As a result, psychiatric disorders might be best thought of as malconnection rather than disconnection syndromes. For example, as described in the previous chapter, a growing body of evidence suggests that certain forms of depression appear to involve alterations in the way circuits in the hippocampus, as well as in the prefrontal cortex and amygdala, adapt to the consequences of long-term elevations of stress hormones. And just as a brain lesion in one area can affect the functions mediated by other regions or systems with which it is connected, so, too, can alterations in the synaptic operation of a region. The only thing one brain area knows about another is the state of its synapses. Change the synapses in one area, and like dominoes in a line, synapses in others will be altered as well.

Most of the time the brain holds the self together pretty well. But when connections change, personality, too, can change. That the self is so fragile an entity is disconcerting. At the same time, if the self can be disassembled by experiences that alter connections, presumably it also can be reassembled by experiences that establish, change, or renew connections. An important challenge for the field of neuroscience is to figure out how to manipulate the brain in a way that patients with mental disorders can, either alone or with the help of a therapist, try to put the self's synapses back together.

SELF-ASSEMBLY

As described in chapter, your brain was assembled during childhood by a combination of genetic and environmental influences. Genes dictated that your brain was a human one and that your synaptic connections, though more similar to those of members of your family than to those of members of other families, were nevertheless distinct. Then, through experiences with the world, your synaptic connections were adjusted (by selection and/or instruction and construction), further distinguishing you from everyone else.

Synaptic connections are adjusted by environmentally driven neural activity in specific neural systems. When these changes occur during early life, they are said to involve developmental plasticity; when they occur later, they are considered as learning. But the line between developmental plasticity and learning is a fine one and perhaps nonexistent. I will therefore ignore this distinction, and plunge right into the question of how synaptic plasticity occurring in multiple neural systems is coordinated in the process of assembling, and maintaining, the self. The manner in which this occurs can, I think, be understood in terms of seven principles.

PRINCIPLE 1 DIFFERENT SYSTEMS EXPERIENCE THE SAME WORLD.

Although the different neural systems have different functions, because they are part of the same brain they

will be involved in encoding the same life events. One system processes the sights, another the sounds, and still another the smells in a given scene. Additional systems will determine whether those sights, sounds, or smells indicate that danger is present, or that there might be something tasty to eat out there. From the point of view of the organism, and the world with which it is interacting, these are not different experiences, but rather different aspects of a single experience. And although each system is plastic, and can thus learn and store information, each is learning and storing information about the same experience. Just as people living in different towns of a country who never meet face-to-face can share a culture because they have similar environmental influences (similar climate, similar geography, similar myths and legends, similar political histories, similar current political situation, similar social institutions), within the brain, a kind of shared culture develops between the various systems because they are exposed to similar environmental circumstances.

To make this principle more explicit, I've put it in pictorial form in figure 11.2. Part A shows a hypothetical brain with three neural systems. Each system receives inputs from the outside world. In this model brain, the systems do not communicate directly with one another. But because they have the same inputs, they come to know exactly the same things about the world. They are thus completely redundant processors.

Part B adds a bit of complexity. Now the three systems are distinct and nonredundant—each encodes the world differently (as represented by the different fills in the shapes). These can be thought of, for example, as different sensory systems—visual, auditory, and olfactory, if you like. In spite of the fact that each processes qualitatively different information, because they experience the same events, there is still a strong degree of overlap in what they represent about the world. The sounds, sights, and smells that occur during an experience are different pieces of information from the point of view of each of the three systems, but are all part of the same experience from the point of view of the brain and person.

Now, imagine that our hypothetical brain visits three different places where it has three different experiences. Each experience will be encoded in parallel by the three systems as sights, sounds, and smells. Contrast this to what would take place if, during each of the three experiences, only one of the neural systems was active, and the one that was active differed in each situation.

In this case, although the organism and its brain had three unique experiences, each system had only one. As a result, this brain has information about one environment in the form of sights, the second in sounds, and the third in smells, but because no parallel encoding took place, there is no shared information about the three experiences across the three systems.

Normally, this condition does not occur. The various systems of your brain share the same experiences. They encode them differently, but they encode the same external events. They will not always focus on the same details, and each may not always participate in every experience. But to the extent that a neural system encodes an experience, it is likely that some other systems of the brain are encoding the same experience. As a result of parallel encoding by, and parallel plasticity within, neural systems, a shared culture develops and persists among the systems, even if they never communicate directly.

PRINCIPLE 2 SYNCHRONY COORDINATES PARALLEL PLASTICITY.

In real brains neural networks do not exist in isolation. They communicate with other networks by way of synaptic transmission. For example, in order to see an apple, instead of a roundish, reddish blob, the various features of the stimulus, each processed by different visual subsystems, have to be integrated. As we saw in chapter 7, the problem of understanding the manner in which this occurs is called *the binding problem*. One popular solution to this problem is based on the notion of neuronal synchrony. Synchronous (simultaneous) firing, and thus binding, has been proposed as an explanation of consciousness (chap. 7), but our interest here is more in the ability of synchronous firing between cells in different interconnected regions to coordinate plasticity across the regions.

Wolf Singer is one of the major proponents of synchrony as a means of integrating plasticity across regions,

especially within the visual system. In brief, his basic idea is that information processing across different interconnected regions is coordinated when cells in the individual regions fire action potentials synchronously—that is, at the same time. Form and color, for example, are brought together for an immediately present object by the fact that the cells processing the particular form and particular color are active at the same time. By way of the synaptic interconnections between cells in the color and form regions, Hebbian plasticity occurs (since the cells will be activating each other at the same time they are being activated by the external visual stimulus). Hebbian plasticity thus binds simultaneously active cells together so that the next time the same or similar stimulus occurs, the same cells and connections will be activated. That synchronous firing can lead to Hebbian plasticity (as opposed to conscious perception) is incontrovertible. In fact, simultaneous (or near simultaneous) activation of inputs is what accounts for Hebbian plasticity (chap. 6).

Unfortunately, little is known about whether changes of this type actually take place between networks in the brain (as opposed to within individual networks). However, recent studies using computer simulations have begun to explore how plasticity occurring in individual systems alters information processing in interconnected systems. This work provides a foundation for pursuing similar studies in real brains.

Integration across brain regions (binding) is usually discussed in the context of perception. However, understanding long-distance communication in the brain is also important for memory; emotion, motivation, and other systems. Exploration of interactions across systems is going to be especially crucial as we try to come to grips with the relation of the self to the brain. A ridiculously simple example is the fact that our perception of an apple is not just based on the integration of the shape and form and other visual features of the object, but also on the integration of these features with information stored in memory about the object and our experiences with it, and its significance for us at the moment, in the past, and in the future.

PRINCIPLE 3 PARALLEL PLASTICITY IS ALSO COORDINATED BY MODULATORY SYSTEMS.

Parallel processing in different brain systems is further coordinated by modulators. As we've seen, these are released throughout the brain in the presence of significant stimuli, including novel, unexpected, or painful stimuli, or stimuli that otherwise signal emotional arousal. In the last chapter, we examined the role of one class of modulators, the monoamines, in mental disorders. Here we examine their role in normal brain function. These are two sides of the same coin.

In the middle of the twentieth century, researchers discovered a region of the brain stem that was required for alertness and arousal. Damage to this region put animals and people into a comatose state. Stimulation of the same region could awaken an animal from a deep sleep, and if an animal was already awake, alertness and attention were enhanced by the stimulation. This area came to be called the “reticular activating system” or the “reticular formation.” Subsequently, it was discovered that arousal functions were largely accounted for not by a single integrated system in the brain stem but by the activities of several different groups of neurons in the vicinity, with each group having a unique chemical signature (different neurotransmitter molecules are present in the different groups). The chemicals in question are all amines, and specifically include the monoamines—dopamine, norepinephrine, epinephrine, serotonin—and acetylcholine.

The cells that produce modulators are located primarily in the brain stem, but their axons are distributed throughout the brain. Consequently, when these cells are activated, many brain areas are affected. The widespread action of modulators makes them especially useful in broadcasting that something significant has happened, but they are less suited to identifying exactly what it is that's happening. The modulatory system functions somewhat like an alarm sounded by the firehouse in the center of a small town. The alarm is very effective in alerting all the town's firemen to the fire, and in summoning them to the station, but it doesn't tell them whose house is on fire. This they have to learn by other means, just as brain areas have to determine

precisely what it is that's causing the arousal by other means.

The main job of modulators is to regulate neurotransmission between neurons, but they don't work at all the many synapses that they bathe. They are mainly effective in modulating transmission at synapses that are already active when the modulator arrives.

Because modulatory systems are activated during significant experiences, modulators can selectively facilitate transmission at the synapses actively processing information about such experiences across widely distributed neural systems. Emotional or otherwise significant experiences are the ones we tend to form memories about, and as described in chapter 8, it is well established that modulators like norepinephrine are involved in the enhancement of memory that occurs during emotional events. Norepinephrine has also been implicated in the induction of long-term potentiation (LTP), which, as we've seen, is a laboratory procedure for studying synaptic plasticity. Thus, LTP is facilitated when this chemical is present, and disrupted when it is absent. So not only can modulators produce a momentary facilitation of transmission in circuits actively involved in processing significant events, they can also promote synaptic plasticity, and thus learning and memory, in those circuits. Recall that one of the most prominent current theories supporting monoamine treatment of mental disorders is that these drugs make more serotonin and/or norepinephrine available at synapses and thereby trigger intracellular molecular cascades that promote synaptic plasticity. Modulation of plasticity across brain systems is thus important in both normal and pathological mental conditions.

One of the most important features of modulators is that, once released, they have a prolonged action, at least with respect to transmitters like glutamate or GABA. The primary action of glutamate or GABA is typically concluded within a matter of milliseconds, whereas modulators can have effects that last for seconds. Given that not all brain systems operate at exactly the same rate (some involve more distant pathways with more connections, and, in general, the more complex the process, the longer it takes to occur because more connections are involved), the slow recovery time of modulators allows them to affect a wide range of processes, from the earliest and simplest in an episode to the last and most complex, promoting learning independent of information extracted during different components of an experience.

Although different systems learn about different aspects of an experience, the widespread action of modulators increases the likelihood that when something significant takes place, plasticity will occur in parallel at active synapses in all these systems. As a result, the learning of multiple elements of an experience

(its sights, sounds, and smells, its emotional and motivational significance, its movement patterns, and so on) is facilitated, allowing the whole experience to be stored at once, albeit across multiple systems. Included, of course, would be systems that process information both implicitly and explicitly, but more about that later.

Not all modulators have the same effect (some inhibit rather than enhance plasticity), and the same modulator can have different effects depending on the particular postsynaptic receptor with which it interacts. Moreover, the same interaction between a modulator and its receptor can be different depending on the other cells in the circuit. For example, when serotonin interacts with one of its receptors, it produces inhibition, and when it interacts with another receptor, it produces excitation. But the net effect (that is, whether excitation or inhibition is the outcome) will depend on the kind of neurons on which the serotonin receptors are located. Serotonin, for instance, inhibits the activity of amygdala projection cells (excitatory cells that transfer information from one region of the amygdala to another). But it achieves this inhibition by way of an excitatory serotonin receptor located on GABA cells. Thus, serotonin excites GABA inhibitory cells, and these inhibit amygdala projection cells. While the outcome of the interaction between serotonin and its receptor is excitatory, the outcome for the overall circuit is inhibition. Much more work is needed to better understand the contribution of modulators to routine transmission in specific circuits, and to the induction of plasticity (learning) and its maintenance (memory). But the gaps in our knowledge do not take away from the established role of some modulators in enhancing plasticity, and thus potentially serving as regulators of plasticity across neural systems.

PRINCIPLE 4 CONVERGENCE ZONES INTEGRATE PARALLEL PLASTICITY.

In the examples cited so far, brain systems learn in parallel. While parallel learning is surely an important part of the complex process by which the self is assembled, parallel learning, on its own (even when buttressed by synchrony and modulatory chemicals), is not sufficient to account for the coherent personality of a human being.

Another important mechanism in self-assembly, especially in humans and other primates, is the existence of convergence zones, regions where information from diverse systems can be integrated. Figure 11.5 shows two independent processing units, the outputs of which meet up in a third, a convergence zone. Many kinds of animals have multiple independent learning systems that can be coerced into learning simultaneously by modulatory chemicals and synchronous firing, but only some animals have convergence zones in their cortex. The cognitive sophistication of a mammalian species, in fact, is nicely predicted by the extent of convergence that occurs in its cortex—more is present in humans than in monkeys, for example, and more in monkeys

than in rats. When plasticity occurs simultaneously in two regions that feed into a convergence zone, plasticity is also likely to occur in the convergence zone since it will be the recipient of the high level of activity that occurs when plasticity is being established in the individual regions. Obviously, synchrony and modulation also influence convergence zones, further increasing their potential to integrate information across systems.

Convergence takes place within systems before it takes place between systems. This can be illustrated by considering object recognition in the “what” stream of the visual cortex (chap. 7). The object recognition system, like other cortical processing systems, is hierarchically organized. The later stages depend on the earlier ones, and information representation becomes increasingly complex as the stages are passed through. In the earliest stage, for example, each cell primarily responds to the orientation of a contour or edge of a small part of the stimulus. Across many cells, all the contours that make up the shape of the stimulus are represented. In the next stage, cells receive convergent inputs from the earlier-stage cells. As a result of getting inputs from cells that represent different parts of that object, each cell in the second stage can represent a larger part of the object. This sort of convergence continues through the hierarchy until at the final stage individual cells represent much of the entire object. These latter cells were once affectionately known as “grandmother cells,” since they were supposed to be able to receive all the information needed to represent a stimulus as complex as your grandmother’s face. While it is no longer believed that grandmother cells exist, many scientists do believe that small sets of synaptically connected cells, called ensembles, receive convergent inputs from lower levels in their processing hierarchy, and represent faces, complex scenes, and other objects of perception. This difference is sometimes described as one between “pontifical” cells, which would make final decisions alone about the way things are, and “cardinal” cells, which do things as a small group. Although single cells have been shown to have remarkable capacities, most researchers accept that ensembles rather than single cells underlie mental and behavioral functions.

Once convergence is completed within systems, it begins to occur across systems. In 1970, E. G. Jones and T. P. S. Powell published a landmark study in which they identified several regions in the monkey cortex that receive convergent inputs from the last processing stage of two or more sensory systems in the cortex. Some of the key convergence zones identified were the posterior parietal cortex, the parahippocampal region and areas of the prefrontal cortex. The ability of these regions to integrate diverse kinds of information explains why they are involved in the most sophisticated cognitive functions of the brain. As we’ve seen, areas of the prefrontal cortex are involved in working memory functions, which underlie many aspects of thinking, planning, and decision-making. The posterior parietal area plays an important role in the cognitive control of movement in space in nonhuman primates, and in humans is crucially involved in language comprehension in the left hemisphere and spatial cognition in the right. Rhinal cortical areas are part of the medial temporal lobe memory system. They establish critical links between sensory areas of the cortex and the hippocampus, and thus provide the hippocampus with the raw materials needed to form relations between external stimuli in the process of establishing long-term explicit memories. The hippocampus, too, is a convergence zone; rather than integrating inputs from different sensory systems per se,

receives inputs from other convergence zones, and is thus something of a super-convergence zone.

In convergence zones like the hippocampus, it is possible for completely independent sensory representations to be synthesized into memory representations that transcend the individual systems involved in the initial processing. Thus, while different systems may form independent memories of separate aspects of an experience, memories formed in, or by way of, a convergence zone are multifaceted—they include information extracted from different systems. Such memories reflect the whole experience of the organism, rather than bits and pieces of an experience recorded by other systems. But because the bits and pieces are the raw materials, there is a kind of unity of experience between the memory established by a convergence zone and by its lower connections. And because the hippocampus and other convergence zones receive inputs from modulatory systems, during significant states of arousal plasticity in these networks is coordinated with the plasticity occurring in other systems throughout the brain.

One important consequence of this arrangement is that though memories are formed by systems that function both implicitly and explicitly during significant experiences, the memories are coordinated to some degree>” That is, the elements you are able to consciously remember about an experience overlap often with some of the elements that were also being separately stored implicitly in other systems. Convergence zones such as those in the medial temporal lobe make possible the creation of consciously accessible memories that integrate elements being encoded separately and implicitly in the other systems. But, remember, while the medial temporal lobe system forms memories in a way that allows them to be consciously accessible, these memories only enter consciousness when they are placed in working memory. And once in working memory, memories and thoughts can, as we see next, influence activity back down the processing hierarchy.

PRINCIPLE 5 DOWNWARDLY MOBILE THOUGHTS COORDINATE PARALLEL PLASTICITY.

So far, I've placed much of the burden of assembling the self on processes that work more or less automatically, from the bottom up. But this is only part of the story. Convergent representations built from the bottom up are also used to direct activity back down the processing hierarchies. Thoughts and memories placed in working memory, for example, can influence what we attend to, the way we see things, and the way we act. These executive control functions of working memory, which were discussed in chapter 7, are possible because the prefrontal cortex, like other convergence zones, reciprocates projections. That is, connections are sent back to the regions that provide the convergent inputs. By pulling the right strings (activating the right axons), working memory can direct traffic in the areas with which it is connected, enhancing the processing of stimuli that are relevant to the task on which it is engaged and suppressing the processing of other stimuli>”

The process by which a thought can cause the brain to issue certain orders is known as downward causation. We prove that downward causation exists every time we carry out an intention. Downward causation is only mind-boggling if you believe that thoughts are one phenomenon and brain activities another. It's still a difficult problem even if you view thinking as a form of brain activity, but the nature of the solution is far more obvious in this case.

If a thought is embodied as a pattern of synaptic transmission within a network of brain cells, as must be the case, then it stands to reason that the brain activity that is a thought can influence activity in other brain systems involved in perception, motivation, movement, and the like. But there's one more connection to make. If a thought is a pattern of neural activity in a network, not only can it cause another network to be active, it can also cause another network to change, to be plastic.

All that is required to induce plasticity at a synapse is the right kind of synaptic activity. If cells processing sensory events can undergo plasticity as a result of the kind of activity those events trigger in sensory systems, then why can't cells processing a thought change the connections of the cells with which they communicate? Obviously,

they do; we simply need to learn more about precisely how this happens.

The downward mobility of thought provides a powerful means by which parallel plasticity in neural systems is coordinated. The more elaborate the convergence zones present in a species, the more elaborate will be the cognitive capacity of the species and the more sophisticated will be the ability of information convergence to coordinate plasticity in that species. With thoughts empowered this way, we can begin to see how the way we think about ourselves can have powerful influences on the way we are, and who we become. One's self-image is self-perpetuating.

PRINCIPLE 6 EMOTIONAL STATES MONOPOLIZE BRAIN RESOURCES

Emotions, too, play a key role in organizing brain activity. We've already had a hint of this when we discussed the function of modulators in coordinating parallel plasticity, since emotional stimuli are some of the most potent activators of modulatory systems. But the influence of emotions is much broader than simply activating modulatory systems, and can be illustrated by reconsidering the various ways that the amygdala affects other brain systems when it detects danger.

In the presence of a threatening stimulus, the amygdala sends direct feedback by way of neural connections to sensory areas of the cortex, encouraging these areas to stay focused on those aspects of the stimulus world that are critical. Amygdala feedback also reaches other cortical areas engaged in thinking and explicit memory formation, encouraging them to think certain thoughts and to form certain memories about the current situation. In addition, the amygdala sends connections to arousal networks, causing them to release their modulatory chemicals throughout the brain. The synapses that are actively involved in processing the external world, in thinking about the world, in forming memories about it, and in receiving the amygdala's feedback will thus be enhanced. In addition, plasticity will be facilitated at these active synapses. And interconnections between active cells in different regions and systems that fire synchronously will be linked by the plasticity that is induced. At the same time, bodily responses controlled by the amygdala will be expressed, and these will provide additional feedback to the brain, not only in the form of bodily sensations that are part of the "felt" response of the emotion, but also in the form of hormones that further affect synaptic activity, over a longer time scale than even modulators. The net result is that emotional arousal penetrates the brain widely, and perpetuates itself.

As we've seen, the brain has a number of emotion systems, including networks involved in identifying sexual partners and food sources, as well as detecting and defending against danger. When one of these is active, the others tend to be inhibited. For example, other things being equal, animals will tend to spend time in areas where they are safe. So when it comes time to search for food, their fear of certain places, like wide-open spaces or places where they've previously encountered a predator, has to be overcome if that's where food is likely to be found. The hungrier the animal is, the more it will tolerate fear and anxiety and take risks to obtain nourishment. Similarly, both eating and sexual arousal are decreased by activation of systems involved in fear and stress. But, once aroused, sexual desire can over ride many other brain systems—people risk all sorts of adverse consequences for an adulterous fling. Not only does the arousal of an emotional state bring many of the brain's cognitive resources to bear on that state, it also shuts down other emotion systems. As a result, learning is coordinated across systems in a very specific manner, ensuring that the learning that does occur is relevant to the current emotional situation.

Because emotion systems coordinate learning, the broader the range of emotions that a child experiences the broader will be the emotional range of the self that develops. This is why childhood abuse is so devastating. If a significant proportion of the early emotional experiences one has are due to activation of the fear system rather than positive systems, then the characteristic personality that begins to build up from the parallel learning processes coordinated by the emotional state is one characterized by negativity and hopelessness rather than affection and optimism.

The wide influence of emotional arousal results in many brain systems being activated simultaneously, many more than if one is engaged in quiet cognitive activity, like lying back musing about something, or even when

vigorously thinking about the solution to a problem. And because more brain systems are typically active during emotional than during nonemotional states, and the intensity of arousal is greater, the opportunity for coordinated learning across brain systems is greater during emotional states. By coordinating parallel plasticity throughout the brain, emotional states promote the development and unification of the self.

PRINCIPLE 7 IMPLICIT AND EXPLICIT ASPECTS OF THE SELF OVERLAP, BUT NOT COMPLETELY.

In spite of the multiplicity of checks and balances that help keep the various systems of the brain on one track, learning about the same experiences and about the same things in each experience, the job performed is not always performed perfectly. Sometimes, the things learned explicitly are not the things that were focused on by the implicit systems, especially emotional systems—recall the ability of the amygdala to learn independent of the cortex (chap. 8). Although there are probably many reasons why this is so, the most obvious one is that there is an imperfect set of connections between cognitive and emotional systems in the current stage of evolution of the human brain. This state of affairs is part of the price we pay for having newly evolved cognitive capacities that are not yet fully integrated into our brains. Although this is also a problem for other primates, it is particularly acute for humans,

since the brain of our species, especially our cortex, was extensively rewired in the process of acquiring natural language functions.

Language both required additional cognitive capacities and made new ones possible, and these changes took space and connections to achieve. The space problem was solved, as we saw earlier, by moving some things around in existing cortical space, and also by adding more space. But the connection problem was only partially solved. The part that was solved, connectivity within cortical processing networks, made the enhanced cognitive capacities of the hominid brain possible. But the part that hasn't been fully solved is connectivity between cognitive systems and other parts of the mental trilogy—emotional and motivational systems. This is why a brilliant mathematician or artist, or a successful entrepreneur, can like anyone else fall victim to sexual seduction, road rage, or jealousy, or be a child abuser or rapist, or can have crippling depression or anxiety. Our brain has not evolved to the point where the new systems that make complex thinking possible can easily control the old systems that give rise to our base needs and motives, and emotional reactions. This doesn't mean that we're simply victims of our brains and should just give in to our urges. It means that downward causation is sometimes hard work. *Doing* the right thing doesn't always flow naturally from *knowing* what the right thing to do is.

In the end, then, the self is maintained by systems that function both explicitly and implicitly. Through explicit systems, we try to willfully dictate who we are, and how we will behave. But we are only partially effective in doing so, since we have imperfect conscious access to emotional systems, which play such a crucial role in coordinating learning by other systems. In spite of their importance, though, emotion systems are not always active and have only episodic influence on what other brain systems learn and store. Furthermore, because there are multiple independent emotion systems, the episodic influence of any one system is itself but a component of the total impact of emotions on self-development.

YOU ARE YOUR SYNAPSES

Synaptic connections hold the self together in most of us most of the time. Sometimes, though, thoughts, emotions, and motivations come uncoupled. If the mental trilogy breaks down, the self is likely to begin to disintegrate

and mental health to deteriorate. When thoughts are radically dissociated from emotions and motivations, as in schizophrenia, personality can, in fact, change drastically. When emotions run wild, as in anxiety disorders or depression, a person is no longer the person he or she once was. And when motivations are subjugated by drug addiction, the emotional and intellectual aspects of life suffer.

That the self is synaptic can be a curse—it doesn't take much to break it apart. But it is also a blessing, as there are always new connections waiting to be made. You are your synapses. They are who you are.

THE END.

THIS HAS BEEN. THE SYNAPTIC SELF. BY JOSEPH LEDOUX.