NEW YORK TIMES BESTSELLER

"A readily engaging, literate, and well-researched glimpse into the great mystery of how we make up our minds." —SAN FRANCISCO CHRONICLE

How We Decide

JONAH LEHRER

AUTHOR OF *Proust Was a Neuroscientist*

2

The Predictions of Dopamine

In the early morning hours of February 24, 1991, the First and Second Marine divisions rolled north across the desert of Saudi Arabia. As they approached the unmarked border with Kuwait—the landscape was just an expanse of barren sand—the troops accelerated their pace. These Marines were the first Coalition forces to enter the country since it had been invaded by Iraq, more than eight months earlier. The outcome of Operation Desert Storm depended on their success. The Marines needed to liberate Kuwait, and they needed to do it in fewer than one hundred hours. If the Marines failed to overtake the Iraqi army quickly, they faced the prospect of urban warfare. The Iraqis were threatening to retreat into the streets of Kuwait City, and if that happened, the ground war could drag on for months.

The Marines expected heavy resistance. The Iraqis had fortified many of their military positions inside Kuwait, concentrating their forces near the Al Wafrah oil field along the Saudi Arabian border. They had draped a line of explosive mines across the desert. To make matters even more difficult, these Iraqi units had largely been spared the brutal air war. Because the Coalition forces were determined to minimize collateral damage and civilian casualties, bombing runs inside the occupied country were sharply restricted. Unlike the Republican Guard troops stationed in southern Iraq, a military force that had been decimated by thirty-seven days of intense bombing, these Marines were about to encounter an enemy at full strength. Central Command (CENTCOM) estimated that during the invasion of Kuwait, each Marine division would suffer approximately a thousand casualties, or between 5 and 10 percent of its total troop strength.

To support this high-stakes mission, a fleet of Coalition battleships and destroyers was positioned fewer than twenty miles off the Kuwaiti coast. This was a risky strategic move; although the big naval guns provided crucial air cover for the ground attack of Kuwait, they were also well within range of Iraqi missiles. On the morning of the Marine invasion, the American and British ships in the Persian Gulf were put on the highest possible alert. They were told to expect hostile fire.

The first twenty-four hours of the ground war exceeded even CENTCOM's high expectations. After successfully breaching the perimeter of mines and barbed wire put down by the Iraqis, the Marine division managed to penetrate deep into central Kuwait. Unlike the Soviet T-72 tanks used by the Iraqi army, the American M1 Abrams tanks were equipped with GPS units and thermal sights, allowing the Marines to engage the enemy in the pitch-black night. After a brigade of Marines reached the outskirts of Kuwait City, they made an abrupt turn to the east and began the task of securing the coastline. Just before dawn on February 25, ten Marine helicopters, along with an amphibious landing ship, conducted a feint attack on a military base near the Kuwaiti port of Ash Shuaybah. The attack was supported by a barrage of artillery rounds from the offshore battleships. The Coalition forces weren't interested in capturing the port; they just wanted to "neutralize" it, to make sure it didn't pose a danger to the offshore convoy.

That same morning, while Ash Shuaybah was being attacked, Lieutenant Commander Michael Riley was monitoring the radar screens onboard the HMS *Gloucester,* a British destroyer stationed about fifteen miles from the port. The ship was responsible for protecting the Allied fleet, which meant that Riley had to monitor all of the airspace surrounding the naval convoy. Since the start of the air war, the radar crews had maintained an exhausting schedule. They were on duty for six hours, then they had six hours to sleep and eat, and after that brief respite, they headed back to the claustrophobic radar room. By the time the ground invasion began, the men were showing signs of fatigue. They had bloodshot eyes and needed constant infusions of caffeine.

Riley had been on duty since midnight. At 5:01 in the morning, just as the Allied ships began shelling Ash Shuaybah, he noticed a radar blip off the Kuwaiti coast. A quick calculation of its trajectory had it heading straight for the convoy. Although Riley had been staring at similar-looking blips all night long, there was something about this radar trace that immediately made him suspicious. He couldn't explain why, but the blinking green dot on the screen filled him with fear; his pulse started to race and his hands became clammy. He continued to observe the incoming blip for another forty seconds as it slowly honed in on the USS *Missouri,* an American battleship. With each sweep of the radar, the blip grew closer. It was approaching the American ship at more than 550 miles per hour. If Riley was going to shoot down the target—if he was going to act on his fear—then he needed to respond right away. If that blip was a missile and Riley didn't move immediately, it would be too late. Hundreds of sailors would die. The USS *Missouri* would be sunk. And Riley would have stood by and watched it happen.

But Riley had a problem. The radar blip was located in airspace that was frequently traveled by American A-6 fighter jets, which the U.S. Navy was using to deliver laser-guided bombs to support the Marine ground invasion. After completing their sorties, the planes flew down the Kuwait coast, turned east toward the convoy, and landed on their aircraft carriers. Over the last few weeks, Riley had watched dozens of A-6s fly a route nearly identical to the one being followed by this unidentified radar blip. The blip was also traveling at the same speed as the fighter jets and had a similar surface area. It looked exactly like an A-6 on the radar screen.

To make matters even more complicated, the A-6 pilots had gotten into the bad habit of turning off their electronic identification on their return flights. This identification system allowed Coalition forces to recognize their own, but it also made the planes more vulnerable to Iraqi antiaircraft missiles. Not surprisingly, the pilots opted for the cloak of silence over Iraqi-controlled airspace. As a result, the radar crew onboard the HMS *Gloucester* had no way of contacting this radar blip.

There was one last way for radar crews to distinguish between an incoming missile and a friendly aircraft: they could determine the altitude of the blip. The A-6 generally flew at around three thousand feet, while a Silkworm missile flew at one thousand feet. However, the type of radar that Riley was using didn't provide him with any altitude information. If he wanted to know the height of a specific object, he had to use a specialized radar system known as the 909, which conducted sweeps in horizontal bands. Unfortunately, the 909 radar operator had entered an incorrect tracking number shortly after the blip appeared, which meant that Riley had no way of knowing the altitude of the flying object. Although he'd now been staring at the radar blip for almost a minute, its identity remained a befuddling mystery.

The target was moving fast. The time for deliberation was over. Riley issued the order to fire; two Sea Dart surface-to-air missiles were launched into the sky. Seconds passed. Riley nervously stared at the radar screen, watching his missiles race toward the object at speeds approaching Mach 1. The blinking green blips appeared to be drawn to the target, like iron filings to a magnet. Riley waited for the interception.

The explosion echoed over the ocean. All of the blips immediately disappeared from the radar screen. Whatever had been flying toward the USS *Missouri* helplessly fell into the sea, just seven hundred yards in front of the American battleship. A few moments later, the captain of the HMS *Gloucester* entered the radar room. "Whose bird is it?" he asked Riley, wanting to know who was responsible for destroying the still unidentified target. "It was ours, sir," Riley responded. The captain asked Riley how he could be sure he'd fired at an Iraqi missile and not at an American fighter jet. Riley said he just knew.

THE NEXT FOUR HOURS were the longest ones of Riley's life. If he had shot down an A-6, then he had killed two innocent pilots. His career was over. He might even be court-martialed. Riley immediately went back to review the radar tapes, looking for any scrap of evidence suggesting that the blip really was an Iraqi missile. But even when he had the luxury of time and analysis, Riley still couldn't definitively identify the target; the tapes were completely ambiguous. The mood on the HMS *Gloucester* quickly grew somber. Investigative teams were sent out to view the wreckage still floating on the ocean surface. An immediate inventory of all Coalition planes in the area was conducted.

The captain of the HMS *Gloucester* heard the news first. He walked over to Riley's bunk, where Riley was trying, in vain, to get some sleep. The results of the investigation were in: the radar blip was a Silkworm missile, not an American fighter jet. Riley had single-handedly saved a battleship.

Of course, it's possible that Riley had just gotten lucky. After the war was over, British naval officers carefully analyzed the sequence of events preceding Riley's decision to fire the Sea Dart missiles. They concluded that based on the radar tapes, it was impossible to distinguish between the Silkworm and a friendly A-6. Although Riley had made the correct decision, he could have just as easily been shooting down an American fighter jet. His risky gamble had paid off, but it had still been a gamble.

That, at least, was the official version of events until the summer of 1993, when Gary Klein started to investigate the Silkworm affair. A cognitive psychologist who consults for the Marine Corps, Klein was informed that nobody could explain how the radar blip had been identified as a hostile missile. Even Riley didn't know why he'd considered that early-morning blip so dangerous. He assumed, like everybody else, that he'd just gotten lucky.

Klein was intrigued. He had spent the last few decades studying decision-making in high-pressure situations, and he knew that intuition could often be astonishingly insightful, even if the origin of those insights was obscure. He was determined to find the source of Riley's fear, to figure out why this particular blip had felt so scary. So he went back to the radar tapes.

He soon realized that Riley had gotten used to seeing a very consistent blip pattern when the A-6s returned from their bombing sorties. Because Riley's naval radar could pick up signals only over water—after a signal went "wet feet"—he was accustomed to seeing the fighter jets right as they flew off the Kuwaiti coast. The planes typically became visible after a single radar sweep.

Klein analyzed the radar tapes from the predawn missile attack. He replayed those fateful forty seconds over and over again, searching for any differences between Riley's experience of the A-6s returning from their sorties and his experience of the Silkworm blip.

That's when Klein suddenly saw the discrepancy. It was subtle, but crystal clear. He could finally explain Riley's intuitive insight.

The secret was the timing. Unlike the A-6, the Silkworm didn't appear off the coast right away. Because it traveled at such a low altitude, nearly two thousand feet below an A-6's, the signal of the missile was initially masked by ground interference. As a result, it wasn't visible until the *third* radar sweep, which was eight seconds after an A-6 would have appeared. Riley was unconsciously evaluating the altitude of the blip, even if he didn't know he was doing it.

This is why Riley got the chills when he stared at the Iraqi missile on his radar screen. There was something strange about this radar blip. It didn't feel like an A-6. Although Riley couldn't explain why he felt so scared, he knew that something scary was happening. This blip needed to be shot down.

1

The question still remains: how did Riley's emotions manage to distinguish between these two seemingly identical radar blips? What was happening inside his brain when he first saw the Silkworm missile, three sweeps off the Kuwaiti coast? Where did his fear come from? The answer lies in a single molecule, called dopamine, that brain cells use to communicate with one another. When Riley stared at the radar screen, it was most likely his dopamine neurons that told him he was looking at a missile and not an A-6 fighter jet.

The importance of dopamine was discovered by accident. In 1954, James Olds and Peter Milner, two neuroscientists at McGill University, decided to implant an electrode deep into the center of a rat's brain. The precise placement of the electrode was largely happenstance; at the time, the geography of the mind remained a mystery. But Olds and Milner got lucky. They inserted the needle right next to the nucleus accumbens (NAcc), a part of the brain that generates pleasurable feelings. Whenever you eat a piece of chocolate cake, or listen to a favorite pop song, or watch your favorite team win the World Series, it is your NAcc that helps you feel so happy.

But Olds and Milner quickly discovered that too much pleasure can be fatal. They placed the electrodes in several rodents' brains and then ran a small current into each wire, making the NAccs continually excited. The scientists noticed that the rodents lost interest in everything. They stopped eating and drinking. All courtship behavior ceased. The rats would just huddle in the corners of their cages, transfixed by their bliss. Within days, all of the animals had perished. They died of thirst.

It took several decades of painstaking research, but neuroscientists eventually discovered that the rats had been suffering from an excess of dopamine. The stimulation of the NAcc triggered a massive release of the neurotransmitter, which overwhelmed the rodents with ecstasy. In humans, addictive drugs work the same way: a crack addict who has just gotten a fix is no different than a rat in an electrical rapture. The brains of both creatures have been blinded by pleasure. This, then, became the dopaminergic cliché; it was the chemical explanation for sex, drugs, and rock and roll.

But happiness isn't the only feeling that dopamine produces. Scientists now know that this neurotransmitter helps to regulate *all* of our emotions, from the first stirrings of love to the most visceral forms of disgust. It is the common neural currency of the mind, the molecule that helps us decide among alternatives. By looking at how dopamine works inside the brain, we can see why feelings are capable of providing deep insights. While Plato disparaged emotions as irrational and untrustworthy—the wild horses of the soul—they actually reflect an enormous amount of invisible analysis.

Much of our understanding of the dopamine system comes from the pioneering research of Wolfram Schultz, a neuroscientist at Cambridge University. He likes to compare dopamine neurons (those neurons that use dopamine to communicate) to the photoreceptors on the retina, which detect the rays of light entering the eye. Just as the process of sight starts with the retina, so the process of decision-making begins with the fluctuations of dopamine.

As a medical student in the early 1970s, Schultz grew interested in the neurotransmitter because of its role in triggering the paralyzing symptoms of Parkinson's disease. He recorded from cells in the monkey brain, hoping to find which cells were involved in controlling the body's movements. But he couldn't find anything. "It was a classic case of experimental failure," he says. "I was a very frustrated scientist." But after years of searching, Schultz noticed something odd about these dopamine neurons: they began to fire just before the monkey was given a reward, such as a pellet of food or a bit of banana. (The rewards were used to get the monkeys to move.) "At first I thought it was unlikely that an individual cell could represent anything so complicated as food," Schultz says. "It just seemed like too much information for one neuron."

After hundreds of experimental trials, Schultz began to believe his own data; he realized he had accidentally found the reward mechanism at work in the primate brain. In the mid-1980s, after publishing a series of landmark papers, Schultz set out to decipher this reward circuitry. How exactly did a single cell manage to represent a reward? And why did it fire *before* a reward was given?

The Schultz experiments followed a simple protocol: he sounded a loud tone, waited for a few seconds, and then squirted some drops of apple juice into the mouth of a monkey. While the experiment was unfolding, Schultz was probing the monkey brain with a needle that monitored the electrical activity inside individual cells. At first, the dopamine neurons fired only when the juice was delivered. The cells were responding to the actual reward. However, once the animal learned that the tone preceded the arrival of juice—this required only a few trials—the same neurons began firing at the sound of the tone instead of at the sweet reward. Schultz called these cells "prediction neurons,"

since they were more concerned with *predicting* rewards than actually receiving them. (This process can be indefinitely extended: the dopamine neurons can be made to respond to a light that precedes the tone that precedes the juice, and so on.) Once this simple pattern was learned, the monkey's dopamine neurons became exquisitely sensitive to variations on it. If the cellular predictions proved correct, and the reward arrived right on time, then the primate experienced a brief surge of dopamine, the pleasure of being right. However, if the pattern was violated—if the tone was played but the juice never arrived—then the monkey's dopamine neurons decreased their firing rate. This is known as the prediction-error signal. The monkey felt upset because its predictions of juice were wrong.

What's interesting about this system is that it's all about *expectation.* Dopamine neurons constantly generate patterns based on experience: if this, then that. They learn that the tone predicts the juice, or that the light predicts the tone that predicts the juice. The cacophony of reality is distilled into models of correlation that allow the brain to anticipate what will happen next. As a result, the monkeys quickly learn when to expect their sweet reward.

After refining this set of cellular forecasts, the brain compares these predictions to what actually happens. Once the monkey is taught to expect juice after a certain sequence of events, its dopamine cells carefully monitor the situation. If everything goes according to plan, its dopamine neurons secrete a little burst of enjoyment. The monkey is happy. But if these expectations aren't met—if the monkey doesn't get the promised juice—the dopamine cells go on strike. They instantly send out a signal announcing their mistake and stop releasing dopamine.

The brain is designed to amplify the shock of these mistaken predictions. Whenever it experiences something unexpected—like a radar blip that doesn't fit the usual pattern, or a drop of juice that doesn't arrive—the cortex immediately takes notice.

Within milliseconds, the activity of the brain cells has been inflated into a powerful emotion. Nothing focuses the mind like surprise.

This fast cellular process begins in a tiny area in the center of the brain that is dense with dopamine neurons. Neuroscientists have known for several years that this region, the anterior cingulate cortex (ACC), is involved in the detection of errors. Whenever the dopamine neurons make a mistaken prediction—when they expect juice but don't get it—the brain generates a unique electrical signal, known as error-related negativity. The signal emanates from the ACC, so many neuroscientists refer to this area as the "oh, shit!" circuit.

The importance of the ACC is revealed by the layout of the brain. Like the orbitofrontal cortex, the ACC helps control the conversation between what we know and what we feel. It sits at the crucial intersection between these two different ways of thinking. On the one hand, the ACC is closely connected to the thalamus, a brain area that helps direct conscious attention. This means that if the ACC is startled by some stimulus—like the bang of a gunshot it didn't expect—it can immediately focus on the relevant sensation. It forces the individual to notice the unexpected event.

While the ACC is alerting the consciousness, it's also sending signals to the hypothalamus, which regulates crucial aspects of bodily function. When the ACC is worried about some anomaly—for instance, an errant blip on a radar screen—that worry is immediately translated into a somatic signal as the muscles prepare for action. Within seconds, heart rate increases, and adrenaline pours into the bloodstream. These fleshly feelings compel us to respond to the situation *right away*. A racing pulse and sweaty palms are the brain's way of saying that there's no time to waste. This prediction error is urgent.

But the ACC doesn't just monitor erroneous predictions. It

also helps remember what the dopamine cells have just learned, so that expectations can be quickly adjusted in light of new events. It internalizes the lessons of real life, making sure that neural patterns are completely up to date. If it was predicted that juice would arrive after the tone, but the juice never arrived, then the ACC makes sure that future predictions are revised. The short-term feeling is translated into a long-term lesson. Even if the monkey is unaware of what, exactly, the ACC has memorized, the next time it's waiting for a squirt of juice, its brain cells are prepared. They know exactly when the reward will arrive.

This is an essential aspect of decision-making. If we can't incorporate the lessons of the past into our future decisions, then we're destined to endlessly repeat our mistakes. When the ACC is surgically removed from the monkey brain, the behavior of the primate becomes erratic and ineffective. The monkeys can no longer predict rewards or make sense of their surroundings. Researchers at Oxford performed an elegant experiment that made this deficit clear. A monkey clutched a joystick that moved in two different directions: it could be either lifted or turned. At any given moment, only one of the movements would trigger a reward (a pellet of food). To make things more interesting, the scientists switched the direction that would be rewarded every twenty-five trials. If the monkey had previously gotten in the habit of lifting the joystick in order to get a food pellet, it now had to shift its strategy.

So what did the monkeys do? Animals with intact ACCs had no problem with the task. As soon as they stopped receiving rewards for lifting the joystick, they started turning it in the other direction. The problem was soon solved, and the monkeys continued to receive their pellets of food. However, monkeys that were missing their ACCs demonstrated a telling defect. When they stopped being rewarded for moving the joystick in a certain direction, they were still able (most of the time) to change direction, just like the normal monkeys. However, they were unable to *persist* in this successful strategy and soon went back to moving the joystick in the direction that garnered no reward. They never learned how to consistently find the food, to turn a mistake into an enduring lesson. Because these monkeys couldn't update their cellular predictions, they ended up hopelessly confused by the simple experiment.

People with a genetic mutation that reduces the number of dopamine receptors in the ACC suffer from a similar problem; just like the monkeys, they are less likely to learn from negative reinforcement. This seemingly minor deficit has powerful consequences. For example, studies have found that people carrying this mutation are significantly more likely to become addicted to drugs and alcohol. Because they have difficulty learning from their mistakes, they make the same mistakes over and over. They can't adjust their behavior even when it proves self-destructive.

The ACC has one last crucial feature, which further explains its importance: it is densely populated with a very rare type of cell known as a spindle neuron. Unlike the rest of our brain cells, which are generally short and bushy, these brain cells are long and slender. They are found only in humans and great apes, which suggests that their evolution was intertwined with higher cognition. Humans have about forty times more spindle cells than any other primate.

The strange form of spindle cells reveals their unique function: their antenna-like bodies are able to convey emotions across the entire brain. After the ACC receives input from a dopamine neuron, spindle cells use their cellular velocity—they transmit electrical signals faster than any other neuron—to make sure that the rest of the cortex is instantly saturated in that specific feeling. The consequence of this is that the minor fluctuations of a single type of neurotransmitter play a huge role in guiding our

actions, telling us how we should feel about what we see. "You're probably 99.9 percent unaware of dopamine release," says Read Montague, a professor of neuroscience at Baylor University. "But you're probably 99.9 percent driven by the information and emotions it conveys to other parts of the brain."

WE CAN NOW begin to understand the surprising wisdom of our emotions. The activity of our dopamine neurons demonstrates that feelings aren't simply reflections of hard-wired animal instincts. Those wild horses aren't acting on a whim. Instead, human emotions are rooted in the predictions of highly flexible brain cells, which are constantly adjusting their connections to reflect reality. Every time you make a mistake or encounter something new, your brain cells are busy changing themselves. Our emotions are deeply empirical.

Look, for example, at Schultz's experiment. When Schultz studied those juice-craving monkeys, he discovered that it took only a few experimental trials before the monkeys' neurons knew exactly when to expect their rewards. The neurons did this by continually incorporating the new information, turning a negative feeling into a teachable moment. If the juice didn't arrive, then the dopamine cells adjusted their expectations. Fool me once, shame on you. Fool me twice, shame on my dopamine neurons.

The same process is constantly at work in the human mind. Motion sickness is largely the result of a dopamine prediction error: there is a conflict between the type of motion being experienced—for instance, the unfamiliar pitch of a boat—and the type of motion *expected* (solid, unmoving ground). The result in this case is nausea and vomiting. But it doesn't take long before the dopamine neurons start to revise their models of motion; this is why seasickness is usually temporary. After a few horrible hours, the dopamine neurons fix their predictions and learn to expect the gentle rocking of the high seas.

When the dopamine system breaks down completely—when neurons are unable to revise their expectations in light of reality—mental illness can result. The roots of schizophrenia remain shrouded in mystery, but one of its causes seems to be an excess of certain types of dopamine receptors. This makes the dopamine system hyperactive and disregulated, which means that the neurons of a schizophrenic are unable to make cogent predictions or correlate their firing with outside events. (Most antipsychotic medications work by reducing the activity of dopamine neurons.) Because schizophrenics cannot detect the patterns that actually exist, they start hallucinating false patterns. This is why schizophrenics become paranoid and experience completely unpredictable shifts in mood. Their emotions have been uncoupled from the events of the real world.

The crippling symptoms of schizophrenia serve to highlight the necessity and precision of dopamine neurons. When these neurons are working properly, they are a crucial source of wisdom. The emotional brain effortlessly figures out what's going on and how to exploit the situation for maximum gain. Every time you experience a feeling of joy or disappointment, fear or happiness, your neurons are busy rewiring themselves, constructing a theory of what sensory cues preceded the emotions. The lesson is then committed to memory, so the next time you make a decision, your brain cells are ready. They have learned how to predict what will happen next.

2

Backgammon is the oldest board game in the world. It was first played in ancient Mesopotamia, starting around 3000 B.C. The

game was a popular diversion in ancient Rome, celebrated by the Persians, and banned by King Louis IX of France for encouraging illicit gambling. In the seventeenth century, Elizabethan courtiers codified the rules of backgammon, and the game has changed little since.

The same can't be said about the *players* of the game. One of the best backgammon players in the world is now a software program. In the early 1990s, Gerald Tesauro, a computer programmer at IBM, began developing a new kind of artificial intelligence (AI). At the time, most AI programs relied on the brute computational power of microchips. This was the approach used by Deep Blue, the powerful set of IBM mainframes that managed to defeat chess grand master Garry Kasparov in 1997. Deep Blue was capable of analyzing more than two hundred million possible chess moves per second, allowing it to consistently select the optimal chess strategy. (Kasparov's brain, on the other hand, evaluated only about five moves per second.) But all this strategic firepower consumed a lot of energy: while playing chess, Deep Blue was a fire hazard and required specialized heat-dissipating equipment so that it didn't burst into flames. Kasparov, meanwhile, barely broke a sweat. That's because the human brain is a model of efficiency: even when it's deep in thought, the cortex consumes less energy than a light bulb.

While the popular press was celebrating Deep Blue's stunning achievement—a machine had outwitted the greatest chess player in the world!—Tesauro was puzzled by its limitations. Here was a machine capable of thinking millions of times faster than its human opponent, and yet it had barely won the match. Tesauro realized that the problem with all conventional AI programs, even brilliant ones like Deep Blue's, was their *rigidity*. Most of Deep Blue's intelligence was derived from other chess grand masters, whose wisdom was painstakingly programmed into the machine. (IBM programmers also studied Kasparov's previous chess matches and engineered the software to exploit his recurring strategic mistakes.) But the machine itself was incapable of learning. Instead, it made decisions by predicting the probable outcomes of several million different chess moves. The move with the highest predicted "value" was what the computer ended up executing. For Deep Blue, the game of chess was just an endless series of math problems.

Of course, this sort of artificial intelligence isn't an accurate model of human cognition. Kasparov managed to compete on the same level as Deep Blue even though his mind had far less computational power. Tesauro's surprising insight was that Kasparov's neurons were effective because they had trained themselves. They had been refined by decades of experience to detect subtle spatial patterns on the chessboard. Unlike Deep Blue, which analyzed *every* possible move, Kasparov was able to instantly winnow his options and focus his mental energies on evaluating only the most useful strategic alternatives.

Tesauro set out to create an AI program that acted like Garry Kasparov. He chose backgammon as his paradigm and named the program TD-Gammon. (The *TD* stands for *temporal difference.*) Deep Blue had been preprogrammed with chess acumen, but Tesauro's software began with absolutely zero knowledge. At first, its backgammon moves were entirely random. It lost every match and made stupid mistakes. But the computer didn't remain a novice for long; TD-Gammon was designed to learn from its own experience. Day and night, the software played backgammon against itself, patiently learning which moves were most effective. After a few hundred thousand games of backgammon, TD-Gammon was able to defeat the best human players in the world.

How did the machine turn itself into an expert? Although the mathematical details of Tesauro's software are numbingly complex, the basic approach is simple.* TD-Gammon generates a set of predictions about how the backgammon game will unfold. Unlike Deep Blue, the computer program doesn't investigate every possible permutation. Instead, it acts like Garry Kasparov and generates its predictions from its previous experiences. The software compares these predictions to what actually happens during the backgammon game. The ensuing discrepancies provide the substance of its education, and the software strives to continually decrease this "error signal." As a result, its predictions constantly increase in accuracy, which means that its strategic decisions get more and more effective and intelligent.

In recent years, the same software strategy has been used to solve all kinds of difficult problems, from programming banks of elevators in skyscrapers to determining the schedules of flights. "Anytime you've got a problem with a seemingly infinite number of possibilities"—the elevators and planes can be arranged in any number of sequences—"these sorts of learning programs can be a crucial guide," says Read Montague. The essential distinction between these reinforcement-learning programs and traditional approaches is that these new programs find the optimal solutions by themselves. Nobody tells the computer how to organize the elevators. Instead, it methodically learns by running trials and focusing on its errors until, after a certain number of

*The TD-learning model used by Tesauro was based on the pioneering work of computer scientists Rich Sutton and Andrew Barto. In the early 1980s, when they were grad students at the University of Massachusetts, Amherst, Sutton and Barto tried to develop a model of artificial intelligence that could learn simple rules and behaviors and apply them to achieve a goal. It was an audacious idea; their academic advisers tried to dissuade them from even trying, but the young scientists were stubborn. "It had always been this kind of untouchable goal in computer science," Sutton says. "Marvin Minsky had done his thesis on reinforcement learning and basically given up. He said it was impossible and left the field. Luckily for us, it wasn't impossible. We knew even simple animals could learn like this—nobody teaches a bird how to find a worm—we just didn't know how."

trials, the elevators are running as efficiently as possible. The seemingly inevitable mistakes have disappeared.

This programming method closely mirrors the activity of dopamine neurons. The brain's cells also measure the mismatch between expectation and outcome. They use their inevitable errors to improve performance; failure is eventually turned into success. Take, for example, an experiment known as the Iowa Gambling Task designed by the neuroscientists Antonio Damasio and Antoine Bechara. The game went as follows: a subject—"the player"—was given four decks of cards, two black and two red, and $2,000 of play money. Each card told the player whether he'd won or lost money. The subject was instructed to turn over a card from one of the four decks and to make as much money as possible.

But the cards weren't distributed at random. The scientists had rigged the game. Two of the decks were full of high-risk cards. These decks had bigger payouts ($100), but also contained extravagant punishments ($1,250). The other two decks, by comparison, were staid and conservative. Although they had smaller payouts ($50), they rarely punished the player. If the gambler drew only from those two decks, he would come out way ahead.

At first, the card-selection process was entirely haphazard. There was no reason to favor any specific deck, and so most people sampled from each pile, searching for the most lucrative cards. On average, people had to turn over about fifty cards before they began to draw solely from the profitable decks. It took about eighty cards before the average experimental subject could explain *why* he or she favored those decks. Logic is slow.

But Damasio wasn't interested in logic; he was interested in emotion. While the gamblers in the experiment were playing the card game, they were hooked up to a machine that measured the electrical conductance of their skin. In general, higher levels of

conductance signal nervousness and anxiety. What the scientists found was that after a player had drawn only ten cards, his hand got "nervous" when it reached for the negative decks. Although the subject still had little inkling of which card piles were the most lucrative, his emotions had developed an accurate sense of fear. The emotions knew which decks were dangerous. The subject's feelings figured out the game first.

Neurologically impaired patients who were unable to experience any emotions at all— usually because of damaged orbitofrontal cortices—proved incapable of selecting the right cards. While most people made substantial amounts of money during the experiment, these purely rational people often went bankrupt and had to take out "loans" from the experimenter. Because these patients were unable to associate the bad decks with negative feelings—their hands never developed the symptoms of nervousness—they continued to draw equally from all four decks. When the mind is denied the emotional sting of losing, it never figures out how to win.

How do emotions become so accurate? How do they identify the lucrative decks so quickly? The answer returns us to dopamine, the molecular source of our feelings. By playing the Iowa Gambling Task with a person undergoing brain surgery for epilepsy—the patient was given local anesthesia but remained awake during the surgery—scientists at the University of Iowa and Caltech were able to watch the learning process unfold in real time. The scientists discovered that human brain cells are programmed just like TD-Gammon: they generate predictions about what will happen and then measure the difference between their expectations and the actual results. In the Iowa Gambling Task experiment, if a cellular prediction proved false—for example, if the player chose the bad deck—then the dopamine neurons immediately stopped firing. The player experienced a negative emotion and learned not to draw from that deck again.

(Disappointment is educational.) However, if the prediction was accurate—if he got rewarded for choosing a lucrative card—then the player felt the pleasure of being correct, and that particular connection was reinforced. As a result, his neurons quickly learned how to make money. They had found the secret to winning the gambling game before the player could understand and explain the solution.

This is a crucial cognitive talent. Dopamine neurons automatically detect the subtle patterns that we would otherwise fail to notice; they assimilate all the data that we can't consciously comprehend. And then, once they come up with a set of refined predictions about how the world works, they translate these predictions into emotions. Let's say, for example, that you're given lots of information about how twenty different stocks have performed over a period of time. (The various share prices are displayed on a ticker tape at the bottom of a television screen, just as they appear on CNBC.) You'll soon discover that you have difficulty remembering all the financial data. If somebody asks you which stocks performed the best, you'll probably be unable to give a good answer. You can't process all the information. However, if you're asked which stocks trigger the best *feelings*—your emotional brain is now being quizzed—you'll suddenly be able to identify the best stocks. According to Tilmann Betsch, the psychologist who performed this clever little experiment, your emotions will "reveal a remarkable degree of sensitivity" to the actual performance of all of the different securities. The investments that rose in value will be associated with the most positive emotions, while the shares that went down in value will trigger a vague sense of unease. These wise yet inexplicable feelings are an essential part of the decision-making process. Even when we think we know nothing, our brains know something. That's what our feelings are trying to tell us.

3

This doesn't mean that people can coast on these cellular emotions. Dopamine neurons need to be continually trained and retrained, or else their predictive accuracy declines. Trusting one's emotions requires constant vigilance; intelligent intuition is the result of deliberate practice. What Cervantes said about proverbs—"They are short sentences drawn from long experience"—also applies to brain cells, but only if we use them properly.

Consider Bill Robertie. He's one of the only people in the world who's a world-class expert in three different games. (Imagine if Bo Jackson had played in the NBA in addition to the NFL and baseball's major leagues . . .) Robertie is a chess master and a former winner of the U.S. speed chess championship. He's a widely respected poker expert and best-selling author of several books on Texas hold'em. However, Robertie is best known for his backgammon skills. He has won the World Championship of Backgammon twice (a feat accomplished by only one other person), and is regularly ranked among the top twenty players in the world. In the early 1990s, when Gerald Tesauro was looking for a backgammon expert to compete against TD-Gammon, he chose Robertie. "He wanted the computer to learn from the best," Robertie says. "And I was the best."

Robertie is now in his early sixties, with a shock of graying hair, lidded eyes, and a pair of thick spectacles. He managed to turn a childhood obsession with chess into a lucrative career. When Robertie talks about games, he still speaks with the boyish enthusiasm of someone who can't quite believe that he gets to play for a living. "The first time I competed against TD-Gammon I was incredibly impressed," Robertie says. "It represented a big improvement over any other computer program I'd ever encountered. But I knew that I was still a better player. The next year, however, was a different story. The computer was now a really formidable opponent. It had learned how to play from playing me."

The software program became a backgammon expert by studying its prediction errors. After making a few million mistakes, TD-Gammon was able to join the shortlist of computers, like Deep Blue, that are able to compete with the best human opponents. However, all of these brilliant machines come with a strict limitation: they can each master only one game. TD-Gammon can't play chess, and Deep Blue can't play backgammon. No computer has been able to master poker.

So how did Robertie get so good at such different games? At first glance, chess, backgammon, and poker seem to rely on very different cognitive skills. That's why most backgammon champions tend to play nothing but backgammon; most chess masters don't bother with card games; and most poker players couldn't tell a Latvian Gambit from a French Defense. And yet, Robertie manages to excel in all three domains. According to Robertie, his success has a simple explanation: "I know how to practice," he says. "I know how to make myself better."

In the early 1970s, when Robertie was still just a chess prodigy—he made a living by winning speed chess tournaments—he stumbled upon backgammon. "Right away, I fell in love with the game," he says. "Plus, there was a lot more money in backgammon than speed chess." Robertie bought a book on backgammon strategy, memorized a few opening moves, and then started to play. And play. And play. "You've got to get obsessed," he says. "You've got to reach the point where you're having dreams about the game."

After a few years of intense practice, Robertie had turned himself into one of the best backgammon players in the world. "I knew I was getting good when I could just glance at a board and know what I should do," Robertie says. "The game started to become very much a matter of aesthetics. My decisions increasingly depended on the look of things, so that I could contemplate a move and then see right away if it made my position look better or worse. You know how an art critic can look at a painting and just know if it's a good painting? I was the same way, only my painting was the backgammon board."

But Robertie didn't become a world champion just by playing a lot of backgammon. "It's not the quantity of practice, it's the *quality,*" he says. According to Robertie, the most effective way to get better is to focus on your mistakes. In other words, you need to consciously consider the errors being internalized by your dopamine neurons. After Robertie plays a chess match, or a poker hand, or a backgammon game, he painstakingly reviews what happened. Every decision is critiqued and analyzed. Should he have sent out his queen sooner? Tried to bluff with a pair of sevens? What if he had consolidated his backgammon blots? Even when Robertie wins—and he almost always wins—he insists on searching for his errors, dissecting those decisions that could have been a little bit better. He knows that self-criticism is the secret to self-improvement; negative feedback is the best kind. "That's one of the things I learned from TD-Gammon," Robertie says. "Here was a computer that did nothing but measure what it got wrong. That's all it did. And it was as good as me."

The physicist Niels Bohr once defined an expert as "a person who has made all the mistakes that can be made in a very narrow field." From the perspective of the brain, Bohr was absolutely right. Expertise is simply the wisdom that emerges from cellular error. Mistakes aren't things to be discouraged. On the contrary, they should be cultivated and carefully investigated.

Carol Dweck, a psychologist at Stanford, has spent decades demonstrating that one of the crucial ingredients of successful education is the ability to learn from mistakes. The same strategy that Robertie uses to excel at games is also an essential pedagogic tool. Unfortunately, children are often taught the exact opposite. Instead of praising kids for trying hard, teachers typically praise them for their innate intelligence (being smart). Dweck has shown that this type of encouragement actually backfires, since it leads students to see mistakes as signs of stupidity and not as the building blocks of knowledge. The regrettable outcome is that kids never learn how to learn.

Dweck's most famous study was conducted in twelve different New York City schools and involved more than four hundred fifth-graders. One at a time, the kids were removed from class and given a relatively easy test consisting of nonverbal puzzles. After the child finished the test, the researchers told the student his or her score and provided a single sentence of praise. Half of the kids were praised for their *intelligence*. "You must be smart at this," the researcher said. The other students were praised for their *effort:* "You must have worked really hard."

The students were then allowed to choose between two different subsequent tests. The first choice was described as a more difficult set of puzzles, but the kids were told that they'd learn a lot from attempting it. The other option was an easy test, similar to the test they'd just taken.

When Dweck was designing the experiment, she'd expected the different forms of praise to have a rather modest effect. After all, it was just one sentence. But it soon became clear that the type of compliment given to the fifth-graders dramatically influenced their choice of tests. Of the group of kids that had been praised for their efforts, 90 percent chose the harder set of puzzles. However, of the kids that were praised for their intelligence, most went for the easier test. "When we praise children for their intelligence," Dweck wrote, "we tell them that this is the name of the game: Look smart, don't risk making mistakes."

Dweck's next set of experiments showed how this fear of failure actually inhibited learning. She gave the same fifth-graders yet another test. This test was designed to be extremely difficult—it was originally written for eighth-graders—but Dweck wanted to see how the kids would respond to the challenge. The students who had been praised for their efforts in the initial test worked hard at figuring out the puzzles. "They got very involved," Dweck says. "Many of them remarked, unprovoked, 'This is my favorite test.'" Kids that had initially been praised for their smarts, on the other hand, were easily discouraged. Their inevitable mistakes were seen as signs of failure: perhaps they really weren't smart after all. After taking this difficult test, the two groups of students had to choose between looking at the exams of kids who did worse than them and looking at the exams of those who did better. Students praised for their intelligence almost always chose to bolster their self-esteem by comparing themselves with students who had performed worse on the test. In contrast, kids praised for their hard work were more interested in the higher-scoring exams. They wanted to understand their mistakes, to learn from their errors, to figure out how to do better.

The final round of tests was the same difficulty level as the initial test. Nevertheless, students who'd been praised for their efforts exhibited significant improvement, raising their average score by 30 percent. Because these kids were willing to challenge themselves, even if it meant failing at first, they ended up performing at a much higher level. This result was even more impressive when compared with students who'd been randomly assigned to the "smart" group; they saw their scores drop by an average of nearly 20 percent. The experience of failure had been so discouraging for the "smart" kids that they actually regressed.

The problem with praising kids for their innate intelligence—the "smart" compliment—is that it misrepresents the neural reality of education. It encourages kids to avoid the most useful kind of learning activities, which is learning from mistakes.

Unless you experience the unpleasant symptoms of being wrong, your brain will never revise its models. Before your neurons can succeed, they must repeatedly fail. There are no shortcuts for this painstaking process.

This insight doesn't apply only to fifth-graders solving puzzles; it applies to everyone. Over time, the brain's flexible cells become the source of expertise. Although we tend to think of experts as being weighed down by information, their intelligence dependent on a vast amount of explicit knowledge, experts are actually profoundly intuitive. When an expert evaluates a situation, he doesn't systematically compare all the available options or consciously analyze the relevant information. He doesn't rely on elaborate spreadsheets or long lists of pros and cons. Instead, the expert naturally depends on the emotions generated by his dopamine neurons. His prediction errors have been translated into useful knowledge, which allows him to tap into a set of accurate feelings he can't begin to explain.

The best experts embrace this intuitive style of thinking. Bill Robertie makes difficult backgammon decisions by just "looking" at the board. Thanks to his rigorous practice techniques, he's confident that his mind has already internalized the ideal moves. Garry Kasparov, the chess grand master, obsessively studied his past matches, looking for the slightest imperfection, but when it came time to play a chess game, he said he played by instinct, "by smell, by feel." After Herb Stein finishes shooting a soap opera episode, he immediately goes home and reviews the rough cut. "I watch the whole thing," Stein says, "and I just take notes. I'm looking really hard for my mistakes. I pretty much always want to find thirty mistakes, thirty things that I could have done better. If I can't find thirty, then I'm not looking hard enough." These mistakes are usually little things, so minor that nobody else would notice. But Stein knows that the only way to get it right the next time is to study what he got wrong this time.

Tom Brady spends hours watching game tape every week, critically looking at each of his passing decisions, but when he's standing in the pocket he knows that he can't hesitate before making a throw. It's not an accident that all of these experts have converged on such a similar method. They have figured out how to take advantage of their mental machinery, to steal as much wisdom as possible from their inevitable errors.

And then there's Lieutenant Commander Michael Riley. Before becoming an officer in the Royal Navy, Riley had spent years learning how to interpret the ambiguous blips on a radar screen. In the Royal Navy, the training process for such warfare specialists revolves around realistic battle simulations so that senior lieutenants like Riley can practice decision-making in its proper context. Officers are able to learn from their mistakes without having to shoot anything down.

During the Persian Gulf War, all of this training paid off. Even though Riley had never seen a Silkworm missile before, his mind had learned how to detect it. Because he had been staring at a radar screen for weeks on end, watching dozens of A-6 jets return from sorties off the Kuwaiti coast, Riley's dopamine neurons started to anticipate a consistent sequence of events. The radar pattern of the American planes had been seared into his brain. But then, in the predawn hours following the ground invasion, Riley saw a radar blip that looked slightly different. When the incoming unidentified blip appeared, it was too far out to sea, three sweeps away from the coast. As a result, a dopamine neuron somewhere in Riley's midbrain was surprised. Here was something that didn't fit the pattern, an error of expectation. The cell instantly responded to the surprising turn of events and altered its rate of firing. This electrical message was passed from neuron to neuron until it reached the ACC. Spindle cells publicized this prediction error all over the brain. Riley's years of naval training were summarized in a single flash of fear.

It was just a feeling, but Riley dared to trust it. "Fire two Sea Darts!" he yelled. The defensive missiles were launched into the sky. The battleship was saved.

SO FAR, we've been exploring the surprising intelligence of our emotions. We've seen how the fluctuations of dopamine are translated into a set of prophetic feelings. But emotions aren't perfect. They are a crucial cognitive tool, but even the most useful tools can't solve every problem. In fact, there are certain conditions that consistently short-circuit the emotional brain, causing people to make bad decisions. The best decision-makers know which situations require *less* intuitive responses, and in the next part of the book, we'll look at what those situations are.

3

Fooled by a Feeling

Ann Klinestiver was working as a high school English teacher in a small town in West Virginia when she was diagnosed with Parkinson's disease. She was only fifty-two years old, but the symptoms were unmistakable. While she was standing at the front of her class trying to teach her students some Shakespeare, her hands started to shake uncontrollably. Then her legs went limp. "I lost control of my own body," she says. "I'd look at my arm, and I'd tell it what to do, but it just wouldn't listen."

Parkinson's is a disease of the dopamine system. It begins when dopamine neurons start to die in a part of the brain that controls the body's movements. Nobody knows why these cells die, but once they are gone, the loss is irrevocable. By the time the symptoms of Parkinson's appear, more than 80 percent of these neurons will be dead.

Ann's neurologist immediately put her on Requip, a drug that imitates the activity of dopamine in the brain. (It's part of a class of drugs called dopamine agonists.) While there are many different treatments for Parkinson's patients, all operate on a simi-