

## THE MIND'S EYE

*To gaze is to think.*

—SALVADOR DALI

Past decades had hula hoops, black-light posters, CB radios, and Rubik's cube. The craze of the 1990s is the autostereogram, also called Magic Eye, Deep Vision, and Superstereogram. These are the computer-generated squiggles that when viewed with crossed eyes or a distant gaze spring into a vivid illusion of three-dimensional, razor-edged objects majestically suspended in space. The fad is now five years old and autostereograms are everywhere, from postcards to Web pages. They have been featured in editorial cartoons, in the *Blondie* comic strip, and in situation comedies like *Seinfeld* and *Ellen*. In one episode, the comedian Ellen DeGeneres belongs to a reading club that has chosen a stereogram book as its weekly selection. Ashamed that she cannot see the illusions, she sets aside an evening to train herself, without success. In desperation she joins a support group for people who cannot "get" stereograms.

Visual illusions fascinated people long before the psychologist Christopher Tyler inadvertently created this sensation in his research on binocular (two-eyed) vision. Simpler illusions made up of parallel lines that seem to converge and congruent lines that look unequal have long appeared in cereal-box reading material, Crackerjack prizes, children's museums, and psychology courses. Their fascination is obvious. "Who are you going to believe, me or your own eyes?" says Groucho Marx to Margaret Dumont, playing on our faith that vision is a certain route to knowledge. As the sayings go: I call them as I see them; Seeing is believing; We have an eyewitness; I saw it with my own eyes. But if a devilish

display can make us see things that aren't there, how can we trust our eyes at other times?

Illusions are no mere curiosities; they set the intellectual agenda for centuries of Western thought. Skeptical philosophy, as old as philosophy itself, impugns our ability to know *anything* by rubbing our faces in illusions: the oar in the water that appears bent, the round tower that from a distance looks flat, the cold finger that perceives tepid water as hot while the hot finger perceives it as cold. Many of the great ideas of the Enlightenment were escape hatches from the depressing conclusions skeptical philosophers drew from illusions. We can know by faith, we can know by science, we can know by reason, we can know that we think and therefore that we are.

Perception scientists take a lighter view. Vision may not work all the time, but we should marvel that it works at all. Most of the time we don't bump into walls, bite into plastic fruit, or fail to recognize our mothers. The robot challenge shows that this is no mean feat. The medieval philosophers were wrong when they thought that objects conveniently spray tiny copies of themselves in all directions and the eye captures a few and grasps their shape directly. We can imagine a science-fiction creature that embraces an object with calipers, prods it with probes and dipsticks, makes rubber molds, drills core samples, and snips off bits for biopsies. But real organisms don't have these luxuries. When they apprehend the world by sight, they have to use the splash of light reflected off its objects, projected as a two-dimensional kaleidoscope of throbbing, heaving streaks on each retina. The brain somehow analyzes the moving collages and arrives at an impressively accurate sense of the objects out there that gave rise to them.

The accuracy is impressive because the problems the brain is solving are literally unsolvable. Recall from Chapter 1 that inverse optics, the deduction of an object's shape and substance from its projection, is an "ill-posed problem," a problem that, as stated, has no unique solution. An elliptical shape on the retina could have come from an oval viewed head-on or a circle viewed at a slant. A patch of gray could have come from a snowball in the shade or a lump of coal in the sun. Vision has evolved to convert these ill-posed problems into solvable ones by adding premises: assumptions about how the world we evolved in is, on average, put together. For example, I will explain how the human visual system "assumes" that matter is cohesive, surfaces are uniformly colored, and objects don't go out of their way to line up in confusing arrangements.

When the current world resembles the average ancestral environment, we see the world as it is. When we land in an exotic world where the assumptions are violated—because of a chain of unlucky coincidences or because a sneaky psychologist concocted the world to violate the assumptions—we fall prey to an illusion. That is why psychologists are obsessed with illusions. They unmask the assumptions that natural selection installed to allow us to solve unsolvable problems and know, much of the time, what is out there.

Perception is the only branch of psychology that has been consistently adaptation-minded, seeing its task as reverse-engineering. The visual system is not there to entertain us with pretty patterns and colors; it is contrived to deliver a sense of the true forms and materials in the world. The selective advantage is obvious: animals that know where the food, the predators, and the cliffs are can put the food in their stomachs, keep themselves out of the stomachs of others, and stay on the right side of the cliff top.

The grandest vision of vision has come from the late artificial intelligence researcher David Marr. Marr was the first to describe vision as solving ill-posed problems by adding assumptions about the world, and was a forceful defender of the computational theory of mind. He also offered the clearest statement of what vision is *for*. Vision, he said, "is a process that produces from images of the external world a description that is useful to the viewer and not cluttered with irrelevant information."

It may seem strange to read that the goal of vision is a "description." After all, we don't walk around muttering a play-by-play narration of everything we see. But Marr was referring not to a publicly spoken description in English but to an internal, abstract one in mentalese. What does it mean to see the world? We *can* describe it in words, of course, but we can also negotiate it, manipulate it physically and mentally, or file it away in memory for future reference. *All* these feats depend on construing the world as real things and stuff, not as the psychedelia of the retinal image. We call a book "rectangular," not "trapezoidal," though it projects a trapezoid on the retina. We mold our fingers into a rectangular (not trapezoidal) posture as we reach for it. We build rectangular (not trapezoidal) shelves to hold it, and we deduce that it can support a broken couch by fitting into the rectangular space beneath it. Somewhere in the mind there must be a mental symbol for "rectangle," delivered by vision but available at once to the rest of the verbal and non-

verbal mind. That mental symbol, and the mental propositions that capture the spatial relations among objects ("book lying face down on shelf near door"), are examples of the "description" that Marr charged vision with computing.

If vision did not deliver a description, every mental faculty—language, walking, grasping, planning, imagining—would need its *own* procedure for deducing that the trapezoid on the retina is a rectangle in the world. That alternative predicts that a person who can call a slanted rectangle a "rectangle" may still have to learn how to hold it as a rectangle, how to predict that it will fit into rectangular spaces, and so on. That seems unlikely. When vision deduces the shape of an object that gave rise to a pattern on the retina, all parts of the mind can exploit the discovery. Though some parts of the visual system siphon off information to motor-control circuits that need to react quickly to moving targets, the system as a whole is not dedicated to any one kind of behavior. It creates a description or representation of the world, couched in objects and 3-D coordinates rather than retinal images, and inscribes it on a blackboard readable by all the mental modules.

This chapter explores how vision turns retinal depictions into mental descriptions. We will work our way up from splashes of light to concepts of objects, and beyond them to a kind of interaction between seeing and thinking known as mental imagery. The repercussions reach to the rest of the psyche. We are primates—highly visual creatures—with minds that evolved around this remarkable sense.

## DEEP EYE

Let's begin with the stereograms. How do they work, and why, for some people, don't they work? Despite all the posters, books, and jigsaw puzzles, I have not seen a single attempt at explaining them to the millions of curious consumers. Understanding stereograms is not only a good way to grasp the workings of perception but it is also a treat for the intellect. Stereograms are yet another example of the marvelous contrivances of natural selection, this one inside our own heads.

Autostereograms exploit not one but four discoveries on how to trick the eye. The first, strange to say, is the picture. We are so jaded by photographs, drawings, television, and movies that we forget that they are a


benign illusion. Smears of ink or flickering phosphor dots can make us laugh, cry, even become sexually aroused. Humans have made pictures for at least thirty thousand years, and contrary to some social-science folklore, the ability to see them as depictions is universal. The psychologist Paul Ekman created a furor in anthropology by showing that isolated New Guinean highlanders could recognize the facial expressions in photographs of Berkeley students. (Emotions, like everything else, were thought to be culturally relative.) Lost in the brouhaha was a more basic discovery: that the New Guineans were seeing things in the photographs at all rather than treating them as blotchy gray paper.

The picture exploits projection, the optical law that makes perception such a hard problem. Vision begins when a photon (unit of light energy) is reflected off a surface and zips along a line through the pupil to stimulate one of the photoreceptors (rods and cones) lining the curved inner surface of the eyeball. The receptor passes a neural signal up to the brain, and the brain's first task is to figure out where in the world that photon came from. Unfortunately, the ray defining the photon's path extends out to infinity, and all the brain knows is that the originating patch lies somewhere along the ray. For all the brain knows, it could be a foot away, a mile away, or many light-years away; information about the third dimension, distance from the eye, has been lost in the process of projection. The ambiguity is multiplied combinatorially by the million other receptors in the retina, each fundamentally confused about how far away its stimulating patch lies. Any retinal image, then, could have been produced by an infinite number of arrangements of three-dimensional surfaces in the world (see the diagram on p. 9).

Of course, we don't *perceive* infinite possibilities; we home in on one, generally close to the correct one. And here is an opening for a crafter of illusions. Arrange some matter so that it projects the same retinal image as an object the brain is biased to recognize, and the brain should have no way of telling the difference. A simple example is the Victorian novelty in which a peephole in a door revealed a sumptuously furnished room, but when the door was opened the room was empty. The sumptuous room was in a dollhouse nailed to the door over the peephole.

The painter-turned-psychologist Adelbert Ames, Jr., made a career out of carpentering even stranger illusory rooms. In one, rods and slabs were suspended from wires higgledy-piggledy throughout the room. But when the room was seen from outside through a peephole in a wall, the rods and slabs lined up into a projection of a kitchen chair. In another


room, the rear wall slanted away from left to right, but it had crazy angles that made its left side just short enough to cancel its expansion in perspective, and its right side just tall enough to cancel its contraction. Through a peephole on the opposite side, the wall projected a rectangle. The visual system hates coincidences: it assumes that a regular image comes from something that really *is* regular and that it doesn't just look that way because of the fortuitous alignment of an irregular shape. Ames *did* align an irregular shape to give a regular image, and he reinforced his cunning trick with crooked windows and floor tiles. When a child stands in the near corner and her mother stands in the far one, the child projects a larger retinal image. The brain takes depth into account when assessing size; that's why a looming toddler never seems to dominate her distant parent in everyday life. But now the viewer's sense of depth is a victim of its distaste for coincidence. Every inch of the wall appears the same distance away, so the retinal images of the bodies are interpreted at face value, and Junior towers over Mom. When they change places by walking along the rear wall, Junior shrinks to lapdog size and Mom becomes Wilt Chamberlain. Ames' room has been built in several museums of science, such as the Exploratorium in San Francisco, and you can see (or be seen in) this astonishing illusion for yourself.



Now, a *picture* is nothing but a more convenient way of arranging matter so that it projects a pattern identical to real objects. The mimicking matter sits on a flat surface, rather than in a dollhouse or suspended by wires, and it is formed by smearing pigments rather than by cutting shapes out of wood. The shapes of the smears can be determined without the twisted ingenuity of an Ames. The trick was stated succinctly by Leonardo da Vinci: "Perspective is nothing else than seeing a place behind a pane of glass, quite transparent, on the surface of which the objects behind the glass are drawn." If the painter sights the scene from

a fixed viewing position and copies the contours faithfully, down to the last hair of the dog, a person who then views the painting from the position of the painter would have his eye impaled by the same sheaf of light rays that the original scene projected. In that part of the visual field the painting and the world would be indistinguishable. Whatever assumptions impel the brain to see the world as the world and not as smeared pigment will impel it to see the *painting* as the world and not as smeared pigment.

What are those assumptions? We'll explore them later, but here is a preview. Surfaces are evenly colored and textured (that is, covered with regular grain, weave, or pockmarking), so a gradual change in the markings on a surface is caused by lighting and perspective. The world often contains parallel, symmetrical, regular, right-angled figures lying on the flat ground, which only *appear* to taper in tandem; the tapering is written off as an effect of perspective. Objects have regular, compact silhouettes, so if Object A has a bite taken out that is filled by Object B, A is behind B; accidents don't happen in which a bulge in B fits flush into the bite in A. You can feel the force of the assumptions in these line drawings, which convey an impression of depth.




In practice, realist painters do not daub paint on windows but use visual images from memory and a host of tricks to accomplish the same thing on a canvas. They use grids made of wire or etched in glass, taut strings running from the scene through pinholes in the canvas to a viewing reticle, the camera obscura, the camera lucida, and now the camera Nikon. And, of course, no painter reproduces every hair of the dog. Brush strokes, the texture of the canvas, and the shape of the frame make a painting depart from the idealization of Leonardo's window. Also, we almost always see a painting from a vantage point different from the one the painter assumed in front of his window, and this makes the sheaf of light rays impaling the eye different from the one the real scene would send out. That is why paintings are only partly illusory: we see what the painting depicts, but we simultaneously see it as a painting, not as reality. The canvas and frame tip us off, and remarkably, we use these very clues about picturehood to ascertain our vantage point relative to the painting

and to compensate for its difference from the painter's. We undo the distortion of the picture as if seeing it from the painter's perspective, and interpret the adjusted shapes correctly. The compensation works only up to a point. When we arrive late to a movie and sit in the front row, the difference between our vantage point and the camera's (analogous to the painter at Leonardo's window) is too much of a stretch, and we see warped actors slithering across a trapezoid.

There is another difference between art and life. The painter had to sight the scene from a single vantage point. The viewer peeps at the world from *two* vantage points: his left eye's and his right eye's. Hold out a finger and remain still while you close one eye, then the other. The finger obscures different parts of the world behind it. The two eyes have slightly different views, a fact of geometry called binocular parallax.

Many kinds of animals have two eyes, and whenever they aim forward, so that their fields overlap (rather than aiming outward for a panoramic view), natural selection must have faced the problem of combining their pictures into a unified image that the rest of the brain can use. That hypothetical image is named after a mythical creature with a single eye in the middle of its forehead: the Cyclops, a member of a race of monocular giants encountered by Odysseus in his travels. The problem in making a cyclopean image is that there is no direct way to overlay the views of the two eyes. Most objects fall on different places in the two images, and the difference depends on how far away they are: the closer the object, the farther apart its facsimiles lie in the two eyes' projections. Imagine looking at an apple on a table, with a lemon behind it and cherries in front.





Your eyes are aimed at the apple, so its image lands on each eye's fovea (the dead center of the retina, where vision is sharpest). The apple is at six o'clock in both retinas. Now look at the projections of the cherries, which are nearer. In the left eye they sit at seven o'clock, but in the right eye they sit at *five* o'clock, not seven. The lemon, which is farther, projects an image at five-thirty in the left eye but at six-thirty in the right eye. Objects closer than the point of fixation wander outward toward the temples; objects that are farther squeeze inward toward the nose.


But the impossibility of a simple overlay presented evolution with an opportunity. With a bit of high school trigonometry, one can *use* the difference in an object's projection in the two eyes, together with the angle formed by the two eyes' gaze and their separation in the skull, to calculate how far away the object is. If natural selection could wire up a neural computer to do the trig, a two-eyed creature could shatter Leonardo's window and sense an object's depth. The mechanism is called stereoscopic vision, stereo for short.

Incredibly, for thousands of years no one noticed. Scientists thought that animals have two eyes for the same reason they have two kidneys: as a by-product of a bilaterally symmetrical body plan, and perhaps so that one could serve as a spare if the other got damaged. The possibility of stereo vision escaped Euclid, Archimedes, and Newton, and even Leonardo did not fully appreciate it. He did notice that the two eyes have different views of a sphere, the left eye seeing slightly farther around it on the left and the right eye seeing farther around it on the right. If only he had used a cube in his example instead of a sphere, he would have noticed that the shapes on the retinas are different. Stereo vision was not discovered until 1838, by Charles Wheatstone, a physicist and inventor after whom the "Wheatstone bridge" electrical circuit is named. Wheatstone wrote:


It will now be obvious why it is impossible for the artist to give a faithful representation of any near solid object, that is, to produce a painting which shall not be distinguished in the mind from the object itself. When the painting and the object are seen with both eyes, in the case of the painting two *similar* pictures are projected on the retinae, in the case of the solid object the two pictures are *dissimilar*; there is therefore an essential difference between the impressions on the organs of sensation in the two cases, and consequently between the perceptions formed in the mind; the painting therefore cannot be confounded with the solid object.

The late discovery of stereo vision is surprising, because it is not hard to notice in everyday experience. Keep one eye closed for a few minutes as you walk around. The world is a flatter place, and you might find yourself grazing doorways and spooning sugar into your lap. Of course, the world does not flatten completely. The brain still has the kinds of information that are present in pictures and television, like tapering, occlusion, placement on the ground, and gradients of texture. Most important, it has motion. As you move around, your vantage point changes continuously, making nearby objects whiz by and farther ones budge more slowly. The brain interprets the flow pattern as a three-dimensional world going by. The perception of structure from optical flow is obvious in *Star Trek*, *Star Wars*, and popular computer screen-savers where white dots fleeing the center of the monitor convey a vivid impression of flying through space (though real stars would be too far away to give that impression to a real-life starfleet crew). All these monocular cues to depth allow people who are blind in one eye to get around pretty well, including the aviator Wiley Post and a wide receiver for the New York Giants football team in the 1970s. The brain is an opportunistic and mathematically adroit consumer of information, and perhaps that is why its use of one cue, binocular disparity, eluded scientists for so long.

Wheatstone proved that the mind turns trigonometry into consciousness when he designed the first fully three-dimensional picture, the stereogram. The idea is simple. Capture a scene using *two* of Leonardo's windows, or, more practically, two cameras, each positioned where one eye would be. Place the right picture in front of a person's right eye and the left picture in front of his left eye. If the brain assumes that the two eyes look at one three-dimensional world, with differences in the views coming from binocular parallax, it should be fooled by the pictures and combine them into a cyclopean image in which objects appear at different depths.




But here Wheatstone ran into a problem, one that still challenges all stereoscopic gadgets. The brain physically adjusts the eyes to the depth of a surface in two ways. First, though I have been describing the pupil as if it were a pinhole, in fact it has a lens to accumulate many rays of light emanating from a point in the world and to focus them all at a point on the retina. The closer the object, the more the rays have to be bent for them to converge to a point rather than to a blurry disk, and the fatter the lens of the eye has to be. Muscles inside the eyeball have to thicken the lens to focus nearby objects and flatten it to focus distant objects.



The squeezing is controlled by the focusing reflex, a feedback loop that adjusts the shape of the lens until the fine detail on the retina is at a maximum. (The circuit is similar to the one used in some autofocus cameras.) Poorly focused movies are annoying to watch because the brain keeps trying to eliminate the blur by accommodating the lens, a futile gesture.

The second physical adjustment is to aim the two eyes, which are about two and a half inches apart, at the same spot in the world. The closer the object, the more the eyes must be crossed.



The eyes are crossed and uncrossed by muscles attached to their sides; the muscles are controlled by a brain circuit that tries to eliminate double images. (Seeing double is often a sign that the brain has been poisoned, suffocated, or bruised.) The circuit is similar to the rangefinders in old cameras, in which a prism superimposes the views from two viewfinder windows and the photographer angles the prism (which is geared to the camera lens) until the images line up. The brain uses the rangefinder principle as another source of information about depth, perhaps an indispensable one. Stereo vision gives information only about *relative* depth—depth in front of or behind the point on which the eyes have converged—and feedback from eyeball direction must be used to anchor a sense of absolute depth.

Now here's the problem for the stereoscope designer. The focusing reflex and the eye-crossing reflex are coupled. If you focus on a nearby point to eliminate blur, the eyes converge; if you focus on a distant one, they become parallel. If you converge your eyes on a nearby point to eliminate double vision, the eyes squeeze the lens to close-up focus; if you diverge your eyes on a distant point, they relax for distant focus. The coupling defeats the most straightforward design for a stereoscope, in which a small picture is placed in front of each eye and both eyes point straight ahead, each at its own picture. Pointing the eyes straight ahead is what you do for distant objects, and it drags the focus of each eye to distance vision, blurring the pictures. Focusing the pictures then brings the eyes together, so the eyes are pointing at the same picture rather than each eye aiming at a different one, and that's no good, either. The eyes bob in and out and the lenses thicken and flatten, but not at the right times. To get a stereoscopic illusion, something has to give.

One solution is to uncouple the responses. Many experimental psychologists have trained themselves like fakirs to wrest control of their reflexes and to "free-fuse" stereograms by an act of will. Some cross their eyes at an imaginary point in front of the picture, so that the left eye is staring at the right picture and vice versa, while they focus each eye on the picture behind the imaginary point. Others lock their eyes straight ahead to infinity while maintaining focus. I once took an afternoon out to train myself to do this after I learned that William James said it was a skill every good psychologist should master. But people with lives cannot be expected to show such dedication.

Wheatstone's invention was a bit ungainly because he faced a second problem: the drawings and daguerreotypes of his age were too big to fit

in front of the eyes without overlapping, and people could not point their eyes outward to gaze at one on each side like fish. So he put one picture off to each side, the two facing each other like bookends, and between them he placed two mirrors glued together like the cover of an open book, each mirror reflecting a picture. He then put a prism in front of each mirror and adjusted them so that the two mirrors appeared to be superimposed. When people looked through the prisms and saw the superimposed reflections of the two pictures, the scene in the pictures leapt into three-dimensionality. The advent of better cameras and smaller film led to a simpler, hand-held design that is still with us. Small pictures—as always, photographed from two vantage points positioned like the eyes—are placed side by side with a perpendicular blinker between them and a glass lens in front of each eye. The glass lens relieves the eye of having to focus its nearby picture, and the eye can relax to its infinity setting. That spreads the eyes so they are pointing straight ahead, one at each picture, and the pictures easily fuse.

The stereoscope became the television of the nineteenth century. Victorian-era families and friends spent cozy hours taking turns to view stereo photographs of Parisian boulevards, Egyptian pyramids, or Niagara Falls. Beautiful wooden stereoscopes and the software for them (cards with side-by-side photographs) are still sold in antique stores to avid collectors. A modern version is the ViewMaster, available at tourist traps the world over: an inexpensive viewer that displays a ring of stereo slides of the local attractions.

A different technique, the anaglyph, overlays the two images on one surface and uses clever tricks so that each eye sees only the image intended for it. A familiar example is the notorious red-and-green cardboard eyeglasses associated with the 3-D movie craze of the early 1950s. The left eye's image is projected in red and the right eye's image is projected in green onto a single white screen. The left eye peers at the screen through a green filter, which makes the white background look green and the green lines intended for the other eye invisible; the red lines intended for the left eye stand out as black. Similarly, the red filter over the right eye makes the background red, the red lines invisible, and the green lines black. Each eye gets its own image, and the Sludge Monsters from Alpha Centauri rise out of the screen. An unfortunate side effect is that when the two eyes see very different patterns like the red and green backgrounds, the brain cannot fuse them. It carves the visual field into a patchwork and seesaws between seeing each patch as green or red, a dis-


concerting effect called binocular rivalry. You can experience a milder case by holding a finger a few inches in front of you with both eyes open gazing into the distance so you get a double image. If you pay attention to one of the double images, you will notice that portions slowly become opaque, dissolve into transparency, fill in again, and so on.

A better kind of anaglyph puts polarizing filters, rather than colored filters, over two projector lenses and in the cardboard glasses. The image intended for the left eye is projected from the left projector in light waves that oscillate in a diagonal plane, like this: /. The light can pass through a filter in front of the left eye which has microscopic slits that are also in that orientation, but cannot pass through a filter in front of the right eye with slits in the opposite orientation, like this: \. Conversely, the filter in front of the right eye allows in only the light coming from the right projector. The superimposed images can be in color, and they do not incite rivalry between the eyes. The technique was used to excellent effect by Alfred Hitchcock in *Dial "M" for Murder* in the scene in which Grace Kelly reaches out for the scissors to stab her would-be strangler. The same cannot be said for the film adaptation of Cole Porter's *Kiss Me Kate*, in which a dancer belts out "Too Darn Hot" on a coffee table while flinging scarves at the camera.

Modern anaglyph glasses have panes made of liquid crystal displays (like the numbers on a digital watch) which act as silent, electrically controlled shutters. At any moment one shutter is transparent and the other is opaque, forcing the eyes to take turns at seeing a computer screen in front of them. The glasses are synchronized with the screen, which shows the left eye's image while the left shutter is open and the right eye's image while the right shutter is open. The views alternate too quickly for the eyes to notice the flicker. The technology is used in some virtual reality displays. But the state of the art in virtual reality is a high-tech version of the Victorian stereoscope. A computer displays each image on a little LCD screen with a lens in front of it, mounted in front of each eye on the inside of a helmet or visor.

These technologies all force the viewer to don or peer through some kind of apparatus. The illusionist's dream is a stereogram that can be seen with the naked eye—an autostereogram.


The principle was discovered a century and a half ago by David Brewster, the Scottish physicist who also studied polarized light and invented the kaleidoscope and the Victorian-era stereoscope. Brewster noticed that the repeating patterns on wallpaper can leap out in depth. Adjacent copies of the pattern, say a flower, can each lure one eye into fixating on it. That can happen because identical flowers are positioned at the same places on the two retinas, so the double image looks like a single image. In fact, like a misbuttoned shirt, a whole parade of double images can falsely mesh into a single image, except for the unpaired members at each end. The brain, seeing no double image, is prematurely satisfied that it has converged the eyes properly, and locks them into the false alignment. This leaves the eyes aimed at an imaginary point behind the wall, and the flowers seem to float in space at that distance. They also seem inflated, because the brain does its trigonometry and calculates how big the flower would have to be at that depth to project its current retinal image.



An easy way to experience the wallpaper effect is to stare at a tile wall a few inches away, too close to focus and converge on comfortably. (Many men rediscover the effect as they stand at a urinal.) The tiles in front of each eye easily fuse, creating the surreal impression of a very large tile wall a great distance away. The wall bows outward, and as the head moves from side to side the wall rocks in the opposite direction. Both would have to happen in the world if the wall were really at that distance while projecting the current retinal image. The brain creates those illusions in its headlong attempt to keep the geometry of the whole hallucination consistent.

Brewster also noticed that any irregularity in the spacing of a pair of copies makes them protrude or recess from the rest. Imagine that the flowers pierced by the lines of sight in the diagram are printed a bit closer to each other. The lines of sight are brought together and cross each other closer to the eyes. The images on the retina will splay out to the temples, and the brain sees the imaginary flower as being nearer. Similarly, if the flowers had been printed a bit farther apart, the lines of sight will cross farther away, and their retinal projections will crowd toward the nose. The brain hallucinates the ghost object at a slightly greater distance.

We have now arrived at a simple kind of "magic eye" illusion, the wall-paper autostereogram. Some of the stereograms in the books and greeting cards show rows of repeating figures—trees, clouds, mountains, people. When you view the stereogram, each tier of objects drifts in or out and lands at its own depth (although in these autostereograms, unlike the squiggly ones, no new shape emerges; we'll come to those soon). Here is an example, designed by Ilavenil Subbiah.



It is like Brewster's wallpaper, but with the unequal separations put in deliberately rather than by a paperhanger's sloppiness. The picture accommodates seven sailboats because they are closely packed, but only five arches because they are spaced farther apart. When you look behind the picture, the sailboats seem closer than the arches because their misbuttoned lines of sight meet in a nearer plane.

If you don't already know how to fuse stereograms, try holding the book right up to your eyes. It is too close to focus; just let your eyes point straight ahead, seeing double. Slowly move the book away while keeping your eyes relaxed and "looking through" the book to an imaginary point



beyond it. (Some people place a pane of glass or a transparency on top of the stereogram, so they can focus on the reflections of distant objects.) You should still be seeing double. The trick is to let one of the double images drift on top of the other, and then to keep them there as if they were magnets. Try to keep the images aligned. The superimposed shapes should gradually come into focus and pop in or out to different depths. As Tyler has noted, stereo vision is like love: if you're not sure, you're not experiencing it.


Some people have better luck holding a finger a few centimeters in front of the stereogram, focusing on the finger, and then removing it while keeping the eyes converged to that depth. With this technique, the false fusion comes from the eyes crossing so that the left eye sights a boat on the right while the right eye sights a boat on the left. Don't worry about what your mother said; your eyes will not freeze into that position forever. Whether you can fuse stereograms with your eyes crossed too much or not enough probably depends on whether you are slightly cross-eyed or wall-eyed to begin with.

With practice, most people can fuse wallpaper autostereograms. They do not need the yogi-like concentration of the psychologists who free-fuse the two-picture stereograms, because they do not have to uncouple their focusing reflex from their convergence reflex to the same degree. Free-fusing a two-picture stereogram requires jamming your eyes far enough apart that each eye remains aimed at one of the pictures. Fusing a wallpaper stereogram requires merely keeping the eyes far enough apart that each eye remains aimed at neighboring clones *inside a single picture*. The clones are close enough together that the convergence angle is not too far out of line from what the focusing reflex wants it to be. It shouldn't be too hard for you to exploit this small wiggle in the mesh between the two reflexes and focus a wee bit closer than your eyes converge. If it is, Ellen DeGeneres may be able to get you into her support group.



The trick behind the wallpaper stereogram—identical drawings luring the eyes into mismatching their views—uncovers a fundamental problem the brain has to solve to see in stereo. Before it can measure the positions of a spot on the two retinas, the brain has to be sure that the

spot on one retina came from the same mark in the world as the spot on the other retina. If the world had only one mark in it, it would be easy. But add a second mark, and their retinal images can be matched in two ways: spot 1 in the left eye with spot 1 in the right eye, and spot 2 in the left eye with spot 2 in the right eye—the correct matchup—or spot 1 in the left eye with spot 2 in the right eye, and spot 2 in the left eye with spot 1 in the right eye—a mismatch that would lead to the hallucination of two ghost marks instead.



Add more marks, and the matching problems multiply. With three marks, there are six ghost matches; with ten marks, ninety; with a hundred marks, almost ten thousand. This "correspondence problem" was noticed in the sixteenth century by the astronomer Johannes Kepler, who thought about how stargazing eyes match up their thousands of white dots and how an object's position in space could be determined from its multiple projections. The wallpaper stereogram works by coaxing the brain to accept a plausible but false solution to the correspondence problem.

Until recently, everyone thought that the brain solved the correspondence problem in everyday scenes by first *recognizing* the objects in each eye and then matching up images of the same object. Lemon in left eye goes with lemon in right eye, cherries in left eye go with cherries in right eye. Stereo vision, guided by the intelligence of the whole person, could head off the mismatches by only joining up points that came from the same kind of object. A typical scene may contain millions of dots, but it will contain far fewer lemons, maybe only one. So if the brain matched whole objects, there would be fewer ways for it to go wrong.


But nature did not opt for that solution. The first hint came from another of Ames' wacky rooms. This time the indefatigable Ames built an ordinary rectangular room but glued leaves on every inch of its floor,

walls, and ceiling. When the room was viewed with one eye through a peephole, it looked like an amorphous sea of green. But when it was viewed with both eyes, it sprang into its correct three-dimensional shape, Ames had built a world that could be seen only by the mythical Cyclopean eye, not by the left eye or the right eye alone. But how could the brain have matched up the two eyes' views if it had to depend on recognizing and linking the objects in each one? The left eye's view was "leaf leaf leaf leaf leaf leaf leaf." The right eye's view was "leaf leaf leaf leaf leaf leaf leaf." The brain was faced with the hardest correspondence problem imaginable. Nonetheless it effortlessly coupled the views and conjured up a cyclopean vision.

The demonstration is not airtight. What if the edges and corners of the room were not perfectly masked by the leaves? Perhaps each eye had a rough inkling of the room's shape, and when the brain fused the two images it became more confident that the inklings were accurate. The airtight proof that the brain can solve the correspondence problem without recognizing objects came from an ingenious early use of computer graphics by the psychologist Bela Julesz. Before he fled Hungary for the United States in 1956, Julesz was a radar engineer with an interest in aerial reconnaissance. Spying from the air uses a clever trick: stereo views penetrate camouflage. A camouflaged object is covered with markings resembling the background it lies on, making the boundary between the object and its background invisible. But as long as the object is not pancake-flat, when it is viewed from *two* vantage points its markings will appear in slightly different positions in the two views, whereas the background markings will not have moved quite as much because they are farther away. The trick in aerial reconnaissance is to photograph the land, let the plane fly a bit, and photograph it again. The pictures are placed side by side and then fed into a hypersensitive detector of disparity in two images: a human being. A person literally looks at the photographs with a stereo viewer, as if he were a giant peering down from the sky with one eye at each position from which the airplane took a picture, and the camouflaged objects pop out in depth. Since a camouflaged object, by definition, is near-invisible in a single view, we have another example of the cyclopean eye seeing what neither real eye can see.

The proof had to come from perfect camouflage, and here Julesz went to the computer. For the left eye's view, he had the computer make a square covered with random dots, like television snow. Julesz then had

the computer make a copy for the right eye, but with one twist: he shifted a patch of dots a bit over to the left, and inserted a new stripe of random dots into the gap at the right so the shifted patch would be perfectly camouflaged. Each picture on its own looked like pepper. But when put in the stereoscope, the patch levitated into the air.



Many authorities on stereo vision at the time refused to believe it because the correspondence problem the brain had to solve was just too hard. They suspected that Julesz had somehow left little cut marks behind in one of the pictures. But of course the computer did no such thing. Anyone who sees a random-dot stereogram is immediately convinced.

All it took for Julesz' occasional collaborator, Christopher Tyler, to invent the magic-eye autostereogram was to combine the wallpaper autostereogram with the random-dot stereogram. The computer generates a vertical stripe of dots and lays copies of it side by side, creating random-dot wallpaper. Say each stripe is ten dots wide, and woinumber the dots from 1 to 10 (using "0" for 10):

```
12345678901234567890i234567890123<67890123456789C1234S67819C
123456789C1234567B901234567890123456789012345678901234567890
12345678901234567S901234567850123456789012345678901234567890
```

—and so on. Any clump of dots—say, "5678"—repeats itself every ten spaces. When the eyes fixate on neighboring stripes, they falsely fuse, just as they do with a wallpaper stereogram, except that the brain is superimposing stretches of random dots rather than flowers. Remember that in a wallpaper stereogram, copies of a pattern that have been squashed closer together will float above the rest because their lines of sight cross closer to the viewer. To make a patch float out of a rragic-eye


\*                      \*

↓                      ↓

[illegible]

You should also see a cutout window lower in the picture. I made it by picking out a rectangular patch and doing the opposite of what I did before: I stuffed an *extra* dot (labeled “X”) next to every copy of dot 4 inside the patch. That pushes the clumps of dots farther apart, so they

repeat themselves every *eleven* spaces. (The stuffed rows, you will notice, are longer than the rest.) Copies that are more widely spaced equals a surface that is more distant. A real random-dot autostereogram, of course, is made of dots, not numbers, so you don't notice the snipped-out or stuffed-in material, and the uneven lines are filled out with extra dots. Here is an example. The fun in viewing a real random-dot autostereogram is that the moment of pop-out surprises the viewer with previously invisible shapes:



When the autostereogram craze hit Japan, it soon developed into an art form. Dots are not necessary; any tapestry of small contours rich enough to fool the brain into locking the eyes on neighboring stripes will do. The first commercial autostereograms used colored squiggles, and the Japanese ones use flowers, ocean waves, and, taking a leaf out of Ames' book, leaves. Thanks to the computer, the shapes don't have to be flat cutouts like in a diorama. By reading in the three-dimensional coordinates of the points on a surface, the computer can shift every dot by a slightly different amount to sculpt the solid shape in cyclopean space, rather than shifting the entire patch rigidly. Smooth, bulbous

shapes materialize, looking as if they are shrink-wrapped in leaves or flowers.

Why did natural selection equip us with true cyclopean vision—an ability to see shapes in stereo that neither eye can see in mono—rather than with a simpler stereo system that would match up the lemons and cherries that are seeable by each eye? Tyler points out that our ancestors really did live in Ames' leaf room. Primates evolved in trees and had to negotiate a network of branches masked by a veil of foliage. The price of failure was a long drop to the forest floor below. Building a stereo computer into these two-eyed creatures must have been irresistible to natural selection, but it could have worked only if the disparities were calculated over thousands of bits of visual texture. Single objects that allow unambiguous matches were just too few and far between.

Julesz points out another advantage of cyclopean vision. Camouflage was discovered by animals long before it was discovered by armies. The earliest primates were similar to today's prosimians, the lemurs and tarsiers of Madagascar, who snatch insects off trees. Many insects hide from predators by freezing, which defeats the hunter's motion detectors, and by camouflage, which defeats its contour detectors. Cyclopean vision is an effective countermeasure, revealing the prey just as aerial reconnaissance reveals tanks and planes. Advances in weaponry spawn arms races in nature no less than in war. Some insects have outwitted their predators' stereo vision by flattening their bodies and lying flush against the background, or by turning into living sculptures of leaves and twigs, a kind of three-dimensional camouflage.

**N**ow does the cyclopean eye work? The correspondence problem—matching up the marks in one eye with their counterparts in the other—is a fearsome chicken-and-egg riddle. You can't measure the stereo disparity of a pair of marks until you have picked a pair of marks to measure. But in a leaf room or a random-dot stereogram, there are thousands of candidates for the matchmaker. If you knew how far away the surface was, you would know where to look on the left retina to find the mate of a mark on the right. But if you knew that, there would be no need to do the stereo computation; you would already have the answer. How does the mind do it?

David Marr noted that built-in assumptions about the world we evolved in can come to the rescue. Among the  $n^2$  possible matches of  $n$  points, not all are likely to have come from this goodly frame, the earth. A well-engineered matcher should consider only the matchups that are physically likely.

First, every mark in the world is anchored to one position on one surface at one time. So a legitimate match must pair up identical points in the two eyes that came from a single splotch in the world. A black dot in one eye should match a black dot in the other, not a white dot, because the matchup has to represent a single position on some surface, and that position cannot be a black splotch and a white splotch at the same time. Conversely, if a black dot does match a black dot, they must come from a single position on some surface in the world. (That is the assumption violated by autostereograms: each of their splotches appears in *several* positions.)

Second, a dot in one eye should be matched with no more than one dot in the other. That means that a line of sight from one eye is assumed to end at a splotch on one and only one surface in the world. At first glance it looks as if the assumption rules out a line of sight passing through a transparent surface to an opaque one, like the bottom of a shallow lake. But the assumption is more subtle; it only rules out the coincidence in which two identical splotches, one on the lake's surface and one on the bottom, line up one behind the other from the left eye's vantage point while both being visible from the right eye's.

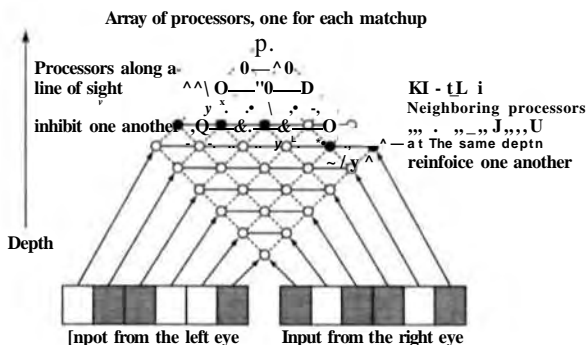
Third, matter is cohesive and smooth. Most of the time a line of sight will end up on a surface in the world that is not drastically closer or farther than the surface hit by the neighboring line of sight. That is, neighboring patches of the world tend to lie on the same smooth surface. Of course, at the boundary of an object the assumption is violated: the edge of the back cover of this book is a couple of feet away from you, but if you glance just to its right you might be looking at the moon a quarter of a million miles away. But boundaries make up a small portion of the visual field (you need much less ink to sketch a line drawing than to color it in), and these exceptions can be tolerated. What the assumption rules out is a world made up of dust storms, swarms of gnats, fine wires, deep crevasses between craggy peaks, beds of nails viewed point-on, and so on.

The assumptions sound reasonable in the abstract, but something still has to *find* the matches that satisfy them. Chicken-and-egg prob-




terns can sometimes be solved with the technique called constraint satisfaction that we met in Chapter 2 when looking at Necker cubes and accented speech. When the parts of a puzzle cannot be solved one at a time, the puzzle-solver can keep in mind several guesses for each one, compare the guesses for the different parts of the puzzle, and see which ones are mutually consistent. A good analogy is working on a crossword puzzle with a pencil and an eraser. Often a clue for a horizontal word is so vague that several words can be penciled in, and a clue for a vertical word is so vague that several words can be penciled in. But if only one of the vertical guesses shares a letter with any of the horizontal guesses, that pair of words is kept and the others are erased. Imagine doing that for all the clues and squares at once and you have the idea of constraint satisfaction. In the case of solving the correspondence problem in stereo vision, the dots are the clues, the matchups and their depths are the guesses, and the three assumptions about the world are like the rules that say that every letter of every word must sit in a box, every box must have a letter in it, and all the sequences of letters must spell out words.

Constraint satisfaction can sometimes be implemented in a constraint network like the one I presented on page 107. Marr and the theoretical neuroscientist Tomaso Poggio designed one for stereo vision. The input units stand for points, such as the black and white squares of a random-dot stereogram. They feed into an array of units that represent all of the  $M \times 72$  possible matchups of a point in the left eye with some other point in the right eye. When one of these units turns on, the network is guessing that there is a splotch at a particular depth in the world (relative to where the eyes have converged). Here is a bird's-eye view of one plane of the network, showing a fraction of the units.



Array of processors, one for each matchup




The model works as follows. A unit turns on only if it gets the same inputs from the two eyes (black or white), embodying the first assumption (each mark anchored to a surface). Because the units are interconnected, the activation of one unit nudges the activations of its neighbors up or down. Units for different matches lying along the same line of sight inhibit one another, embodying the second assumption (no coincidental markings aligned along a line of sight). Units for neighboring points at nearby depths excite one another, embodying the third assumption (matter is cohesive). The activations reverberate around the network, and it eventually stabilizes, with the activated units tracing out a contour in depth. In the diagram, the filled-in units are showing an edge hovering over its background,

The constraint-satisfaction technique, in which thousands of processors make tentative guesses and hash it out among themselves until a global solution emerges, is consistent with the general idea that the brain works with lots of interconnected processors computing in parallel. It captures some of the psychology, too. When viewing a complicated random-dot stereogram, often you don't see the hidden figure erupt instantaneously. A bit of edge might pop out from the pepper, which then lifts up a sheet, which cleans and straightens a Fuzzy border on the other side, and so on until the whole shape coalesces. We experience the solution emerging, but not the struggle of the processors to come up with it. The experience is a good reminder that as we see and think, dozens of iterations of information processing go on beneath the level of consciousness.

The Marr-Poggio model captures the flavor of the brain's computation of stereo vision, but our real circuitry is surely more sophisticated. Experiments have shown that when people are put in artificial worlds that violate assumptions about uniqueness and smoothness, they don't see as badly as the model predicts. The brain must be using additional kinds of information to help solve the matchup problem, for one thing, the world is not made up of random dots. The brain can match up all the little diagonals, T's, zigzags, inkblots, and other jots and tittles in the two eyes' views (which even a random-dot stereogram has in abundance). There are far fewer false matches among jots and tittles than there are among dots, so the number of matches that have to be ruled out is radically shaved.

Another matchmaking trick is to exploit a different geometric consequence of having two eyes, the one noticed by Leonardo: there are parts of an object that one eye can see but that the other eye cannot. Hold a

pen vertically in front of you, with the clip facing away at eleven o'clock. When you close each eye in turn, you will notice that only the left eye can see the clip; it is hidden from the right eye by the rest of the pen. Was natural selection as astute as Leonardo when it designed the brain, letting it use this valuable clue to an object's boundary? Or does the brain ignore the clue, grudgingly chalking up each mismatch as an exception to the cohesive-matter assumption? The psychologists Ken Nakayama and Shinsuke Shimojo have shown that natural selection did not ignore the clue. They created a random-dot stereogram whose depth information lay not in shifted dots but in dots that were visible in one eye's view and absent in the other's. Those dots lay at the corners of an imaginary square, with dots at the top and bottom right corners only in the right eye's picture, and dots in the top and bottom left corners only in the left eye's picture. When people view the stereogram, they see a floating square defined by the four points, showing that the brain indeed interprets features visible to only one eye as coming from an edge in space. Nakayama and the psychologist Barton Anderson suggest that there are neurons that detect these occlusions; they would respond to a pair of marks in one eye, one of which can be matched with a mark in the other eye and the other of which cannot be matched. These 3-D boundary detectors would help a stereo network home in on the outlines of the floating patches.



Stereo vision does not come free with the two eyes; the circuitry has to be wired into the brain. We know this because about two percent of the population can see perfectly well out of each eyeball but not with the cyclopean eye; random-dot stereograms remain flat. Another four percent can see stereo only poorly. An even larger minority has more selective deficits. Some can't see stereo depth behind the point of fixation; others can't see it in front. Whitman Richards, who discovered these forms of stereoblindness, hypothesized that the brain has three pools of neurons that detect differences in the position of a spot in the two eyes. One pool is for pairs of spots that coincide exactly or almost exactly, for fine-grained depth perception at the point of focus. Another is for pairs of spots flanking the nose, for farther objects. A third is for pairs of spots approaching the temples, for nearer objects. Neurons with all these

properties have since been found in the brains of monkeys and cats. The different kinds of stereoblindness appear to be genetically determined, suggesting that each pool of neurons is installed by a different combination of genes.

Stereo vision is not present at birth, and it can be permanently damaged in children or young animals if one of the eyes is temporarily deprived of input by a cataract or a patch. So far, this sounds like the tiresome lesson that stereo vision, like everything else, is a mixture of nature and nurture. But a better way of thinking about it is that the brain has to be assembled, and the assembly requires project scheduling over an extended timetable. The timetable does not care about when the organism is extruded from the womb; the installation sequence can carry on after birth. The process also requires, at critical junctures, the intake of information that the genes cannot predict.

Stereo vision appears abruptly in infants. When newborns are brought into a lab at regular intervals, for week after week they are unimpressed by stereograms, and then suddenly they are captivated. Close to that epochal week, usually around three or four months of age, the babies converge their eyes properly for the first time (for example, they smoothly track a toy brought up to their nose), and they find rivalrous displays—a different pattern in each eye—annoying, whereas before they had found them interesting.

It is not that babies "learn to see in stereo," whatever that would mean. The psychologist Richard Held has a simpler explanation. When infants are born, every neuron in the receiving layer of the visual cortex *adds up* the inputs from corresponding locations in the two eyes rather than keeping them separate. The brain can't tell which eye a given bit of pattern came from, and simply melts one eye's view on top of the other's in a 2-D overlay. Without information about which eye a squiggle came from, stereo vision, convergence, and rivalry are logically impossible. Around the three-month mark each neuron settles on a favorite eye to respond to. The neurons lying one connection downstream can now know when a mark falls on one spot in one eye and on the same spot, or a slightly shifted-over spot, in the other eye—the grist for stereo vision.

In cats and monkeys, whose brains have been studied directly, this is indeed what happens. As soon as the animal's cortex can tell the eyes apart, the animal sees stereograms in depth. That suggests that when the inputs are first tagged "left eye" or "right eye," the circuitry for stereo computation one layer downstream is already installed and functioning.

In monkeys it's all over in two months: by then each neuron has a favorite eye and the baby monkeys see in depth. Compared with other primates, humans are "altricial": babies are born early and helpless, and complete their development outside the womb. Because human infants are born earlier than monkeys in proportion to the length of their childhood, the installation of their binocular circuitry appears at a later age as measured from the date of birth. More generally, when biologists compare the milestones of the maturation of the visual systems of different animals, some born early and helpless, others born late and seeing, they find that the sequence is pretty much the same whether the later steps take place in the womb or in the world.

The emergence of the crucial left-eye and right-eye neurons can be disrupted by experience. When the neurobiologists David Hubel and Torsten Wiesel raised kittens and baby monkeys with one eye covered, the input neurons of the cortex all tuned themselves to the other eye, making the animal functionally blind in the eye that was covered. The damage was permanent, even with only brief deprivation, if the eye was covered in a critical period in the animal's development. In monkeys, the visual system is especially vulnerable during the first two weeks of life, and the vulnerability tapers off during the first year. Covering the eye of an adult monkey, even for four years, does no harm.

At first this all looked like a case of "use it or lose it," but a surprise was in store. When Hubel and Wiesel covered *both* eyes, the brain did not show twice the damage; half the cells showed no damage at all. The havoc in the single-eyepatch experiment came about not because a neuron destined for the covered eye was starved of input but because the input signals from the *uncovered* eye elbowed the covered eye's inputs out of the way. The eyes compete for real estate in the input layer of the cortex. Each neuron begins with a slight bias for one eye or the other, and the input from that eye exaggerates the bias until the neuron responds to it alone. The inputs do not even have to originate in the world; waves of activation from intermediate way-stations, a kind of internally generated test pattern, can do the trick. The developmental saga, though it is sensitive to changes in the animal's experience, is not exactly "learning," in the sense of registering information from the world. Like an architect who hands a rough sketch to a low-level draftsman to straighten out the lines, the genes build eye-specific neurons crudely and then kick off a process that is guaranteed to sharpen them unless a neurobiologist meddles.

Once the brain has segregated the left eye's image from the right eye's, subsequent layers of neurons can compare them for the minute disparities that signal depth. These circuits, too, can be modified by the animal's experience, though again in surprising ways. If an experimenter makes an animal cross-eyed or wall-eyed by cutting one of the eye muscles, the eyes point in different directions and never see the same thing on the two retinas at the same time. Of course, the eyes don't point 180 degrees apart, so in theory the brain could learn to match the nut-of-whack segments that do overlap. But apparently it is not equipped for matches that stretch more than a few degrees across the two eyes; the animal grows up stereoblind, and often functionally blind in one of the two eyes as **well**, a condition called amblyopia. (Amblyopia is sometimes called "lazy eye," but that is misleading. It is the brain, not the eye, that is insensitive, and the insensitivity is caused by the brain actively suppressing one eye's input in a kind of permanent rivalry, not by the brain lazily ignoring it.)

The same thing can happen in children. If one of the eyes is more farsighted than the other, the child habitually strains to focus on nearby objects, and [the reflex that couples focusing and convergence draws that eye inward. The two eyes point in different directions (a condition called strabismus), and their views don't align closely enough for the brain to use the disparity information in them. The child will grow up amblyopic and stereoblind unless early surgery on the eye muscles lines the eyeballs up. Until Hubel and Wiesel discovered these effects in monkeys and Held found similar ones in children, surgery for strabismus was considered cosmetic and done only on school-aged children. But there is a critical period for the proper alignment of two-eye neurons, a bit longer than the one for one-eye neurons but probably fading out near the age of one or two. Surgery after that point is often too late,

Why is there a critical period, as opposed to rigid hard-wiring or life-long openness to experience? In kittens, monkeys, and human babies, the face keeps growing after birth, and the eyes get pushed farther apart. Their relative vantage points change, and the neurons must keep up by retuning the range of intereye disparities they detect. Genes cannot anticipate the degree of spreading of the vantage points, because it depends on other genes, nutrition, and various accidents. So the neurons track the drifting eyes during the window of growth. When the eyes arrive at their grownup separation in the skull, the need disappears, and that is when the critical period ends. Some animals, like rabbits, have

precocious babies whose eyes are set in adult positions within faces that grow very little. (These tend to be prey animals, which don't have the luxury of a long, helpless childhood.) The neurons that receive inputs from the two eyes don't need to retune themselves, and in fact these animals are wired at birth and do without a critical period of sensitivity to the input.

The discoveries about the tunability of binocular vision in different species offers a new way of thinking about learning in general. Learning is often described as indispensable shaper of amorphous brain tissue. Instead it might be an innate adaptation to the project-scheduling demands of a self-assembling animal. The genome builds as much of the animal as it can, and for the parts of the animal that cannot be specified in advance (such as the proper wiring for two eyes that are moving apart at an unpredictable rate), the genome turns on an information-gathering mechanism at the time in development at which it is most needed. In *The Language Instinct* I develop a similar explanation for the critical period for learning language in childhood.


I have led you through magic-eye stereograms not just because it is fun to understand how the magic works. I think stereo vision is one of the glories of nature and a paradigm of how other parts of the mind might work. Stereo vision is information processing that we experience as a particular flavor of consciousness, a connection between mental computation and awareness that is so lawful that computer programmers can manipulate it to enchant millions. It is a module in several senses: it works without the rest of the mind (not needing recognizable objects), the rest of the mind works without it (getting by, if it has to, with other depth analyzers), it imposes particular demands on the wiring of the brain, and it depends on principles specific to its problem (the geometry of binocular parallax). Though stereo vision develops in childhood and is sensitive to experience, it is not insightfully described as "learned" or as "a mixture of nature and nurture"; the development is part of an assembly schedule and the sensitivity to experience is a circumscribed intake of information by a structured system. Stereo vision shows off the engineering acumen of natural selection, exploiting subtle theorems in optics rediscovered millions of years later by the likes of Leonardo da Vinci,



Kepler, Wheatstone, and aerial reconnaissance engineers. It evolved in response to identifiable selection pressures in the ecology of our ancestors. And it solves unsolvable problems by making tacit assumptions about the world that were true when we evolved but are not always true now.

## LIGHTING, SHADING. SHAPING

Stereo vision is part of a crucial early stage of vision that figures out the depths and materials of surfaces, but it is not the only part. Seeing in three dimensions doesn't require two eyes. You can get a rich sense of shape and substance from the meagerest hints in a picture. Look at these drawings, designed by the psychologist Edward Adelson.



The left one appears to be white cardboard with a gray vertical stripe, folded horizontally and lit from above. The right one appears to be white cardboard with a gray horizontal stripe, folded vertically and lit from the side. (If you stare long enough, either might flip in depth, like a Necker cube; let's ignore that for now.) But the ink on the page (and the projection on your retina) is virtually the same in the two pictures. Each is a zigzag tic-tac-toe box with some of the squares shaded in. In both drawings, the corner squares are white, the top and side squares are light gray, and the middle square is a darker gray. Somehow the combination of shading and zigzagging pops them into the third dimension and colorizes each square, but in different ways. The borders labeled "1" are physically the same in the two drawings. But in the left drawing the border looks like a paint boundary—a white stripe next to a gray one—and in the right drawing it looks like a shape-and-shading boundary—a white stripe falling into a shadow on the other side of a fold. The borders labeled "2" are also identical, but you see them in the opposite way: shadow in the