THE FABRIC OF THE COSMOS

SPACE, TIME, AND THE TEXTURE OF REALITY

BRIAN GREENE

PENGUIN BOOKS

This seems like a convincing story. Statistical and probabilistic reasoning has given us the second law of thermodynamics. In turn, the second law has provided us with an intuitive distinction between what we call past and what we call future. It has given us a practical explanation for why things in daily life, things that are typically composed of huge numbers of constituents, start like *this* and end like *that,* while we never see them start like *that* and end like *this.* But over the course of many years—and thanks to important contributions by physicists like Lord Kelvin, Josef Loschmidt, Henri Poincaré, S. H. Burbury, Ernst Zermelo, and Willard Gibbs—Ludwig Boltzmann came to appreciate that the full story of time's arrow is more surprising. Boltzmann realized that although entropy had illuminated important aspects of the puzzle, it had *not* answered the question of why the past and the future seem so different. Instead, entropy had redefined the question in an important way, one that leads to an unexpected conclusion.

Entropy: Past and Future

Earlier, we introduced the dilemma of past versus future by comparing our everyday observations with properties of Newton's laws of classical physics. We emphasized that we continually experience an obvious directionality to the way things unfold in time but the laws themselves treat what we call forward and backward in time on an exactly equal footing. As there is no arrow within the laws of physics that assigns a direction to time, no pointer that declares, "Use these laws in this temporal orientation but not in the reverse," we were led to ask: If the laws underlying experience treat both temporal orientations symmetrically, why are the experiences themselves so temporally lopsided, always happening in one direction but not the other? Where does the observed and experienced directionality of time come from?

In the last section we seemed to have made progress, through the second law of thermodynamics, which apparently singles out the future as the direction in which entropy increases. But on further thought it's not that simple. Notice that in our discussion of entropy and the second law, we did not modify the laws of classical physics in any way. Instead, all we did was use the laws in a "big picture" statistical framework: we ignored fine details (the precise order of *War and Peace*'s unbound pages, the precise locations and velocities of an egg's constituents, the precise locations and velocities of a bottle of Coke's CO_2 molecules) and instead focused our attention on gross, overall features (pages ordered vs. unordered, egg splattered vs. not splattered, gas molecules spread out vs. not spread out). We found that when physical systems are sufficiently complicated (books with many pages, fragile objects that can splatter into many fragments, gas with many molecules), there is a huge difference in entropy between their ordered and disordered configurations. And this means that there is a huge likelihood that the systems will evolve from lower to higher entropy, which is a rough statement of the second law of thermodynamics. But the key fact to notice is that the second law is *derivative:* it is merely a consequence of probabilistic reasoning applied to Newton's laws of motion.

This leads us to a simple but astounding point: *Since Newton's laws of physics have no built-in temporal orientation, all of the reasoning we have used to argue that systems will evolve from lower to higher entropy toward the future works equally well when applied toward the past.* Again, since the underlying laws of physics are time-reversal symmetric, there is no way for them even to distinguish between what we call the past and what we call the future. Just as there are no signposts in the deep darkness of empty space that declare this direction up and that direction down, there is nothing in the laws of classical physics that says this direction is time future and that direction is time past. The laws offer no temporal orientation; it's a distinction to which they are completely insensitive. And since the laws of motion are responsible for how things change—both toward what we call the future and toward what we call the past—the statistical/probabilistic reasoning behind the second law of thermodynamics applies equally well in both temporal directions. Thus, *not only is there an overwhelming probability that the entropy of a physical system will be higher in what we call the future, but there is the same overwhelming probability that it was higher in what we call the past.* We illustrate this in Figure 6.2.

This is *the* key point for all that follows, but it's also deceptively subtle. A common misconception is that if, according to the second law of thermodynamics, entropy increases toward the future, then entropy necessarily *decreases* toward the past. But that's where the subtlety comes in. The second law actually says that if at any given moment of interest, a physical system happens not to possess the maximum possible entropy, it is extraordinarily likely that the physical system will subsequently have *and* previously had more entropy. That's the content of Figure 6.2b. With laws that are blind to the past-versus-future distinction, such time symmetry is inevitable.

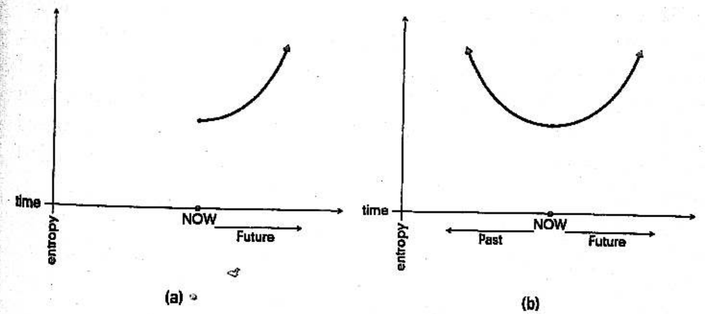

Figure 6.2 **(a)** As it's usually described, the second law of thermodynamics implies that entropy increases toward the future of any given moment. **(b)** Since the known laws of nature treat forward and backward in time identically, the second law actually implies that entropy increases both toward the future and toward the past from any given moment.

That's the essential lesson. It tells us that the entropic arrow of time is *double-headed.* From any specified moment, the arrow of entropy increase points toward the future *and* toward the past. And that makes it decidedly awkward to propose entropy as the explanation of the *one-way* arrow of experiential time.

Think about what the double-headed entropic arrow implies in concrete terms. If it's a warm day and you see partially melted ice cubes in a glass of water, you have full confidence that half an hour later the cubes will be more melted, since the more melted they are, the more entropy they have.[11] But you should have *exactly* the same confidence that half an hour earlier they were also more melted, since *exactly* the same statistical reasoning implies that entropy should increase toward the past. And the same conclusion applies to the countless other examples we encounter every day. Your assuredness that entropy increases toward the future—from partially dispersed gas molecules' further dispersing to partially jumbled page orders' getting more jumbled—should be matched by *exactly* the same assuredness that entropy was also higher in the past.

The troubling thing is that half of these conclusions seem to be flat-out wrong. Entropic reasoning yields accurate and sensible conclusions when applied in one time direction, toward what we call the future, but gives apparently inaccurate and seemingly ridiculous conclusions when applied toward what we call the past. Glasses of water with partially melted ice cubes do not usually start out as glasses of water with no ice cubes in which molecules of water coalesce and cool into chunks of ice, only to start melting once again. Unbound pages of *War and Peace* do not usually start thoroughly out of numerical order and through subsequent tosses get less jumbled, only to start getting more jumbled again. And going back to the kitchen, eggs do not generally start out splattered, and then coalesce into a pristine whole egg, only to splatter some time later.

Or do they?

Following the Math

Centuries of scientific investigations have shown that mathematics provides a powerful and incisive language for analyzing the universe. Indeed, the history of modern science is replete with examples in which the math made predictions that seemed counter to both intuition and experience (that the universe contains black holes, that the universe has anti-matter, that distant particles can be entangled, and so on) but which experiments and observations were ultimately able to confirm. Such developments have impressed themselves profoundly on the culture of theoretical physics. Physicists have come to realize that mathematics, when used with sufficient care, is a proven pathway to truth.

So, when a mathematical analysis of nature's laws shows that entropy should be higher toward the future *and* toward the past of any given moment, physicists don't dismiss it out of hand. Instead, something akin to a physicists' Hippocratic oath impels researchers to maintain a deep and healthy skepticism of the apparent truths of human experience and, with the same skeptical attitude, diligently follow the math and see where it leads. Only then can we properly assess and interpret any remaining mismatch between physical law and common sense.

Toward this end, imagine it's 10:30 p.m. and for the past half hour you've been staring at a glass of ice water (it's a slow night at the bar), watching the cubes slowly melt into small, misshapen forms. You have absolutely no doubt that a half hour earlier the bartender put fully formed ice cubes into the glass; you have no doubt because you trust your memory. And if, by some chance, your confidence regarding what happened during the last half hour should be shaken, you can ask the guy across the way, who was also watching the ice cubes melt (it's a *really* slow night at the bar), or perhaps check the video taken by the bar's surveillance camera, both of which would confirm that your memory is accurate. If you were then to ask yourself what you expect to happen to the ice cubes during the next half hour, you'd probably conclude that they'd continue to melt. And, if you'd gained sufficient familiarity with the concept of entropy, you'd explain your prediction by appealing to the overwhelming likelihood that entropy will increase from what you see, right now at 10:30 p.m., toward the future. All that makes good sense and jibes with our intuition and experience.

But as we've seen, such entropic reasoning—reasoning that simply says things are more likely to be disordered since there are more ways to be disordered, reasoning which is demonstrably powerful at explaining how things unfold toward the future—proclaims that entropy is just as likely to also have been higher in the past. This would mean that the partially melted cubes you see at 10:30 p.m. would actually have been *more* melted at earlier times; it would mean that at 10:00 p.m. they did not begin as solid ice cubes, but, instead, slowly coalesced out of room-temperature water on the way to 10:30 p.m., just as surely as they will slowly melt into room-temperature water on their way to 11:00 p.m.

No doubt, that sounds weird—or perhaps you'd say nutty. To be true, not only would H_2O molecules in a glass of room-temperature water have to coalesce spontaneously into partially formed cubes of ice, but the digital bits in the surveillance camera, as well as the neurons in your brain and those in the brain of the guy across the way, would all need to spontaneously arrange themselves by 10:30 p.m. to attest to there having been a collection of fully formed ice cubes that melted, even though there never was. Yet this bizarre-sounding conclusion is where a faithful application of entropic reasoning—the same reasoning that you embrace without hesitation to explain why the partially melted ice you see at 10:30 p.m. continues to melt toward 11:00 p.m.—leads when applied in the time-symmetric manner dictated by the laws of physics. This is the trouble with having fundamental laws of motion with no inbuilt distinction between past and future, laws whose mathematics treats the future and past of any given moment in exactly the same way.[12]

Rest assured that we will shortly find a way out of the strange place to which an egalitarian use of entropic reasoning has taken us; I'm not going to try to convince you that your memories and records are of a past that never happened (apologies to fans of *The Matrix*). But we will find it very useful to pinpoint precisely the disjuncture between intuition and the mathematical laws. So let's keep following the trail.

A Quagmire

Your intuition balks at a past with higher entropy because, when viewed in the usual forward-time unfolding of events, it would require a spontaneous rise in order: water molecules spontaneously cooling to 0 degrees Celsius and turning into ice, brains spontaneously acquiring memories of things that didn't happen, video cameras spontaneously producing images of things that never were, and so on, all of which seem extraordinarily unlikely—a proposed explanation of the past at which even Oliver Stone would scoff. On this point, the physical laws and the mathematics of entropy agree with your intuition completely. Such a sequence of events, when viewed in the forward time direction from 10 p.m. to 10:30 p.m., goes against the grain of the second law of thermodynamics—it results in a decrease in entropy—and so, although not impossible, it *is* very unlikely.

By contrast, your intuition and experience tell you that a far more likely sequence of events is that ice cubes that were fully formed at 10 p.m. partially melted into what you see in your glass, right now, at 10:30 p.m. But on this point, the physical laws and mathematics of entropy only partly agree with your expectation. Math and intuition concur that *if* there really were fully formed ice cubes at 10 p.m., then the most likely sequence of events would be for them to melt into the partial cubes you see at 10:30 p.m.: the resulting increase in entropy is in line both with the second law of thermodynamics and with experience. But where math and intuition deviate is that our intuition, unlike the math, fails to take account of the likelihood, or lack thereof, of actually having fully formed ice cubes at 10 p.m., *given the one observation we are taking as unassailable, as fully trustworthy, that right now, at 10:30 p.m., you see partially melted cubes.*

This is the pivotal point, so let me explain. The main lesson of the second law of thermodynamics is that physical systems have an overwhelming tendency to be in high-entropy configurations because there are so many ways such states can be realized. And once in such high-entropy states, physical systems have an overwhelming tendency to stay in them. High entropy is the natural state of being. You should never be surprised by or feel the need to explain why any physical system is in a high-entropy state. Such states are the norm. On the contrary, what does need explaining is why any given physical system is in a state of order, a state of low

entropy. These states are not the norm. They can certainly happen. But from the viewpoint of entropy, such ordered states are rare aberrations that cry out for explanation. So the one fact in the episode we are taking as unquestionably true—your observation at 10:30 p.m. of low-entropy partially formed ice cubes—is a fact in need of an explanation.

And from the point of view of probability, it is absurd to explain this low-entropy state by invoking the even *lower*-entropy state, the *even less likely* state, that at 10 p.m. there were *even more ordered, more fully formed* ice cubes being observed in a *more* pristine, *more* ordered environment. Instead, it is enormously more likely that things began in an unsurprising, totally normal, high-entropy state: a glass of uniform liquid water with absolutely no ice. Then, through an unlikely but every-so-often-expectable statistical fluctuation, the glass of water went against the grain of the second law and evolved to a state of lower entropy in which partially formed ice cubes appeared. This evolution, although requiring rare and unfamiliar processes, completely avoids the even lower-entropy, the even less likely, the even more rare state of having *fully* formed ice cubes. At every moment between 10 p.m. and 10:30 p.m., this strange-sounding evolution has *higher* entropy than the normal ice-melting scenario, as you can see in Figure 6.3, and so it realizes the accepted observation at 10:30 p.m. in a way that is *more likely*—hugely more likely—than the scenario in which fully formed ice cubes melt.[13] That is the crux of the matter.*

*Remember, on pages 152–53 we showed the huge difference between the number of ordered and disordered configurations for a mere 693 double-sided sheets of paper. We are now discussing the behavior of roughly 10^{24} H_2O molecules, so the difference between the number of ordered and disordered configurations is breathtakingly monumental. Moreover, the same reasoning holds for all other atoms and molecules within you and within the environment (brains, security cameras, air molecules, and so on). Namely, in the standard explanation in which you can trust your memories, not only would the partially melted ice cubes have begun, at 10 p.m., in a more ordered—less likely—state, but so would everything else: when a video camera records a sequence of events, there is a net increase in entropy (from the heat and noise released by the recording process); similarly, when a brain records a memory, although we understand the microscopic details with less accuracy, there is a net increase in entropy (the brain may gain order but as with any order-producing process, if we take account of heat generated, there is a net increase in entropy). Thus, if we compare the total entropy in the bar between 10 p.m. and 10:30 p.m. in the two scenarios—one in which you trust your memories, and the other in which things spontaneously arrange themselves from an initial state of disorder to be consistent with what you see, now, at 10:30 p.m.—there is an enormous entropy difference. The latter scenario, every step of the way, has *hugely* more entropy than the former scenario, and so, from the standpoint of probability, is hugely more likely.

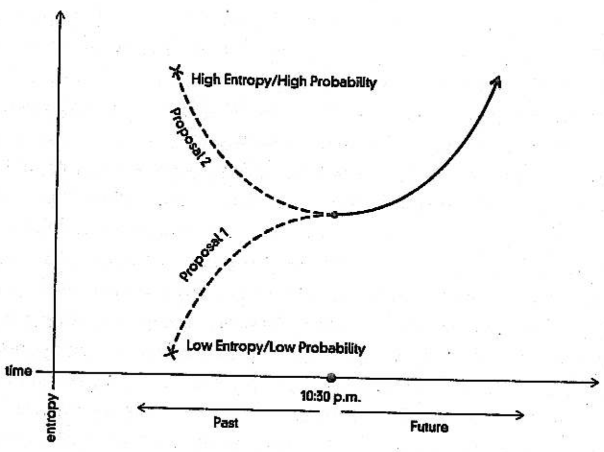

Figure 6.3 A comparison of two proposals for how the ice cubes got to their partially melted state, right now, at 10:30 p.m. Proposal 1 aligns with your memories of melting ice, but requires a comparatively low-entropy starting point at 10:00 p.m. Proposal 2 challenges your memories by describing the partially melted ice you see at 10:30 p.m. as having coalesced out of a glass of water, but starts off in a high-entropy, highly probable configuration of disorder at 10:00 p.m. Every step of the way toward 10:30 p.m., Proposal 2 involves states that are more likely than those in Proposal 1—because, as you can see in the graph, they have higher entropy—and so Proposal 2 is statistically favored.

It was a small step for Boltzmann to realize that the whole of the universe is subject to this same analysis. When you look around the universe right now, what you see reflects a great deal of biological organization, chemical structure, and physical order. Although the universe could be a totally disorganized mess, it's not. Why is this? Where did the order come from? Well, just as with the ice cubes, from the standpoint of probability it is extremely unlikely that the universe we see evolved from an even more ordered—an even less likely—state in the distant past that has slowly unwound to its current form. Rather, because the cosmos has so many constituents, the scales of ordered versus disordered are magnified intensely. And so what's true at the bar is true with a vengeance for the

whole universe: it is *far* more likely—breathtakingly more likely—that the whole universe we now see arose as a statistically rare fluctuation from a normal, unsurprising, high-entropy, completely disordered configuration.

Think of it this way: if you toss a handful of pennies over and over again, sooner or later they will all land heads. If you have nearly the infinite patience needed to throw the jumbled pages of *War and Peace* in the air over and over again, sooner or later they will land in correct numerical order. If you wait with your open bottle of flat Coke, sooner or later the random jostling of the carbon dioxide molecules will cause them to reenter the bottle. And, for Boltzmann's kicker, if the universe waits long enough—for nearly an eternity, perhaps—its usual, high-entropy, highly probable, totally disordered state will, through its own bumping, jostling, and random streaming of particles and radiation, sooner or later just happen to coalesce into the configuration that we all see right now. Our bodies and brains would emerge fully formed from the chaos—stocked with memories, knowledge, and skills—even though the past they seem to reflect would never really have happened. Everything we know about, everything we value, would amount to nothing more than a rare but every-so-often-expectable statistical fluctuation momentarily interrupting a near eternity of disorder. This is schematically illustrated in Figure 6.4.

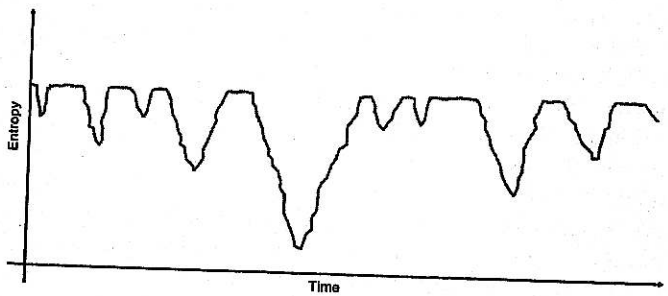

Figure 6.4 A schematic graph of the universe's total entropy through time. The graph shows the universe spending most of its time in a state of total disorder—a state of high entropy—and every so often experiencing fluctuations to states of varying degrees of order, varying states of lower entropy. The greater the entropy dip, the less likely the fluctuation. Significant dips in entropy, to the kind of order in the universe today, are extremely unlikely and would happen very rarely.

Taking a Step Back

When I first encountered this idea many years ago, it was a bit of a shock. Up until that point, I had thought I understood the concept of entropy fairly well, but the fact of the matter was that, following the approach of textbooks I'd studied, I'd only ever considered entropy's implications for the future. And, as we've just seen, while entropy applied toward the future confirms our intuition and experience, entropy applied toward the past just as thoroughly contradicts them. It wasn't quite as bad as suddenly learning that you've been betrayed by a longtime friend, but for me, it was pretty close.

Nevertheless, sometimes it's good not to pass judgment too quickly, and entropy's apparent failure to live up to expectations provides a case in point. As you're probably thinking, the idea that all we're familiar with just popped into existence is as tantalizing as it is hard to swallow. And it's not "merely" that this explanation of the universe challenges the veracity of everything we hold to be real and important. It also leaves critical questions unanswered. For instance, the more ordered the universe is today—the greater the dip in Figure 6.4—the more surprising and unlikely is the statistical aberration required to bring it into existence. So if the universe could have cut any corners, making things look more or less like what we see right now while skimping on the actual amount of order, probabilistic reasoning leads us to believe it would have. But when we examine the universe, there seem to be numerous lost opportunities, since there are many things that are more ordered than they have to be. If Michael Jackson never recorded *Thriller* and the millions of copies of this album now distributed worldwide all got there as part of an aberrant fluctuation toward lower entropy, the aberration would have been far less extreme if only a million or a half-million or just a few albums had formed. If evolution never happened and we humans got here via an aberrant jump toward lower entropy, the aberration would have been far less extreme if there weren't such a consistent and ordered evolutionary fossil record. If the big bang never happened and the more than 100 billion galaxies we now see arose as an aberrant jump toward lower entropy, the aberration would have been less extreme if there were 50 billion, or 5,000, or just a handful, or just one galaxy. And so if the idea that our universe is a statistical fluctuation—a happy fluke—has any validity, one would need to address how and why the universe went so far overboard and achieved a state of *such* low entropy.

Even more pressing, if you truly can't trust memories and records, then you also can't trust the laws of physics. Their validity rests on numerous experiments whose positive outcomes are attested to only by those very same memories and records. So all the reasoning based on the time-reversal symmetry of the accepted laws of physics would be totally thrown into question, thereby undermining our understanding of entropy and the whole basis for the current discussion. By embracing the conclusion that the universe we know is a rare but every-so-often-expectable statistical fluctuation from a configuration of total disorder, we're quickly led into a quagmire in which we lose all understanding, including the very chain of reasoning that led us to consider such an odd explanation in the first place.*

Thus, by suspending disbelief and diligently following the laws of physics and the mathematics of entropy—concepts which in combination tell us that it is overwhelmingly likely that disorder will increase both toward the future *and* toward the past from any given moment—we have gotten ourselves neck deep in quicksand. And while that might not sound pleasant, for two reasons it's a very good thing. First, it shows with precision why mistrust of memories and records—something at which we intuitively scoff—doesn't make sense. Second, by reaching a point where our whole analytical scaffolding is on the verge of collapse, we realize, forcefully, that we *must* have left something crucial out of our reasoning.

Therefore, to avoid the explanatory abyss, we ask ourselves: what new idea or concept, beyond entropy and the time symmetry of nature's laws, do we need in order to go back to trusting our memories and our records—our experience of room-temperature ice cubes melting and not unmelting, of cream and coffee mixing but not unmixing, of eggs splattering but not unsplattering? In short, where are we led if we try to explain

*A closely related point is that should we convince ourselves that the world we see right now just coalesced out of total disorder, the exact same reasoning—invoked anytime later—would require us to abandon our current belief and, instead, attribute the ordered world to a yet more recent fluctuation. Thus, in this way of thinking, every next moment invalidates the beliefs held in each previous moment, a distinctly unconvincing way of explaining the cosmos.

an asymmetric unfolding of events in spacetime, with entropy to our future higher, but entropy to our past *lower*? Is it possible?

It is. But only if things were very special early on.[14]

The Egg, the Chicken, and the Big Bang

To see what this means, let's take the example of a pristine, low-entropy, fully formed egg. How did this low-entropy physical system come into being? Well, putting our trust back in memories and records, we all know the answer. The egg came from a chicken. And that chicken came from an egg, which came from a chicken, which came from an egg, and so on. But, as emphasized most forcefully by the English mathematician Roger Penrose,[15] this chicken-and-egg story actually teaches us something deep and leads somewhere definite.

A chicken, or any living being for that matter, is a physical system of astonishingly high order. Where does this organization come from and how is it sustained? A chicken stays alive, and in particular, stays alive long enough to produce eggs, by eating and breathing. Food and oxygen provide the raw materials from which living beings extract the energy they require. But there is a critical feature of this energy that must be emphasized if we are to really understand what's going on. Over the course of its life, a chicken that stays fit takes in just about as much energy in the form of food as it gives back to the environment, mostly in the form of heat and other waste generated by its metabolic processes and daily activities. If there weren't such a balance of energy-in and energy-out, the chicken would get increasingly hefty.

The essential point, though, is that all forms of energy are not equal. The energy a chicken gives off to the environment in the form of heat is highly disordered—it often results in some air molecules here or there jostling around a touch more quickly than they otherwise would. Such energy has high entropy—it is diffuse and intermingled with the environment—and so cannot easily be harnessed for any useful purpose. To the contrary, the energy the chicken takes in from its feed has low entropy and is readily harnessed for important life-sustaining activities. So the chicken, and every life form in fact, is a conduit for taking in low-entropy energy and giving off high-entropy energy.

This realization pushes the question of where the low entropy of an egg originates one step further back. How is it that the chicken's energy source, the food, has such low entropy? How do we explain this aberrant source of order? If the food is of animal origin, we are led back to the initial question of how animals have such low entropy. But if we follow the food chain, we ultimately come upon animals (like me) that eat only plants. How do plants and their products of fruits and vegetables maintain low entropy? Through photosynthesis, plants use sunlight to separate ambient carbon dioxide into oxygen, which is given back to the environment, and carbon, which the plants use to grow and flourish. So we can trace the low-entropy, nonanimal sources of energy to the sun.

This pushes the question of explaining low entropy another step further back: where did our highly ordered sun come from? The sun formed about 5 billion years ago from an initially diffuse cloud of gas that began to swirl and clump under the mutual gravitational attraction of all its constituents. As the gas cloud got denser, the gravitational pull of one part on another got stronger, causing the cloud to collapse further in on itself. And as gravity squeezed the cloud tighter, it got hotter. Ultimately, it got hot enough to ignite nuclear processes that generated enough outward-flowing radiation to stem further gravitational contraction of the gas. A hot, stable, brightly burning star was born.

So where did the diffuse cloud of gas come from? It likely formed from the remains of older stars that reached the end of their lives, went supernova, and spewed their contents out into space. Where did the diffuse gas responsible for these early stars come from? We believe that the gas was formed in the aftermath of the big bang. Our most refined theories of the origin of the universe—our most refined *cosmological* theories—tell us that by the time the universe was a couple of minutes old, it was filled with a *nearly uniform hot gas* composed of roughly 75 percent hydrogen, 23 percent helium, and small amounts of deuterium and lithium. The essential point is that this gas filling the universe had extraordinarily *low* entropy. The big bang started the universe off in a state of low entropy, and that state appears to be the source of the order we currently see. In other words, *the current order is a cosmological relic.* Let's discuss this important realization in a little more detail.

Entropy and Gravity

Because theory and observation show that within a few minutes after the big bang, primordial gas was uniformly spread throughout the young universe, you might think, given our earlier discussion of the Coke and its carbon dioxide molecules, that the primordial gas was in a high-entropy, disordered state. But this turns out not to be true. Our earlier discussion of entropy completely ignored gravity, a sensible thing to do because gravity hardly plays a role in the behavior of the minimal amount of gas emerging from a bottle of Coke. And with that assumption, we found that uniformly dispersed gas has high entropy. But when gravity matters, the story is very different. Gravity is a universally attractive force; hence, if you have a large enough mass of gas, every region of gas will pull on every other and this will cause the gas to fragment into clumps, somewhat as surface tension causes water on a sheet of wax paper to fragment into droplets. When gravity matters, as it did in the high-density early universe, clumpiness—not uniformity—is the norm; it is the state toward which a gas tends to evolve, as illustrated in Figure 6.5.

Even though the clumps appear to be more ordered than the initially diffuse gas—much as a playroom with toys that are neatly grouped in trunks and bins is more ordered than one in which the toys are uniformly strewn around the floor—in calculating entropy you need to tally up the contributions from *all* sources. For the playroom, the entropy decrease in going from wildly strewn toys to their all being "clumped" in trunks and bins is more than compensated for by the entropy increase from the fat burned and heat generated by the parents who spent hours cleaning and arranging everything. Similarly, for the initially diffuse gas cloud, you find that the entropy decrease through the formation of orderly clumps is more than compensated by the heat generated as the gas compresses, and,

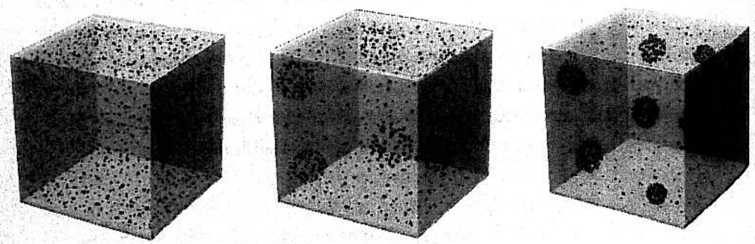

Figure 6.5 For huge volumes of gas, when gravity matters, atoms and molecules evolve from a smooth, evenly spread configuration, into one involving larger and denser clumps.

ultimately, by the enormous amount of heat and light released when nuclear processes begin to take place.

This is an important point that is sometimes overlooked. The overwhelming drive toward disorder does not mean that orderly structures like stars and planets, or orderly life forms like plants and animals, can't form. They can. And they obviously do. What the second law of thermodynamics entails is that in the formation of order there is generally a more-than-compensating generation of disorder. The entropy balance sheet is still in the black even though certain constituents have become more ordered. And of the fundamental forces of nature, gravity is the one that exploits this feature of the entropy tally to the hilt. Because gravity operates across vast distances and is universally attractive, it instigates the formation of the ordered clumps—stars—that give off the light we see in a clear night sky, all in keeping with the net balance of entropy increase.

The more squeezed, dense, and massive the clumps of gas are, the larger the overall entropy. Black holes, the most extreme form of gravitational clumping and squeezing in the universe, take this to the limit. The gravitational pull of a black hole is so strong that nothing, not even light, is able to escape, which explains why black holes are black. Thus, unlike ordinary stars, black holes stubbornly hold on to all the entropy they produce: none of it can escape the black hole's powerful gravitational grip.[16] In fact, as we will discuss in Chapter 16, nothing in the universe contains more disorder—more entropy—than a black hole.* This makes good intuitive sense: high entropy means that many rearrangements of the constituents of an object go unnoticed. Since we can't see inside a black hole, it is impossible for us to detect *any* rearrangement of its constituents—whatever those constituents may be—and hence black holes have maximum entropy. When gravity flexes its muscles to the limit, it becomes the most efficient generator of entropy in the known universe.

We have now come to the place where the buck finally stops. *The ultimate source of order, of low entropy, must be the big bang itself.* In its earliest moments, rather than being filled with gargantuan containers of entropy such as black holes, as we would expect from probabilistic considerations, for some reason the nascent universe was filled with a hot, uniform, gaseous mixture of hydrogen and helium. Although this configu-

*That is, a black hole of a given size contains more entropy than *anything* else of the same size.

ration has high entropy when densities are so low that we can ignore gravity, the situation is otherwise when gravity can't be ignored; then, such a uniform gas has extremely low entropy. In comparison with black holes, the diffuse, nearly uniform gas was in an extraordinarily low-entropy state. Ever since, in accordance with the second law of thermodynamics, the overall entropy of the universe has been gradually getting higher and higher; the overall, net amount of disorder has been gradually increasing. After about a billion years or so, gravity caused the primordial gas to clump, and the clumps ultimately formed stars, galaxies, and some lighter clumps that became planets. At least one such planet had a nearby star that provided a relatively low-entropy source of energy that allowed low-entropy life forms to evolve, and among such life forms there eventually was a chicken that laid an egg that found its way to your kitchen counter, and much to your chagrin that egg continued on the relentless trajectory to a higher entropic state by rolling off the counter and splattering on the floor. The egg splatters rather than unsplatters because it is carrying forward the drive toward higher entropy that was initiated by the extraordinarily low entropy state with which the universe began. Incredible order at the beginning is what started it all off, and we have been living through the gradual unfolding toward higher disorder ever since.

This is the stunning connection we've been leading up to for the entire chapter. *A splattering egg tells us something deep about the big bang.* It tells us that the big bang gave rise to an extraordinarily ordered nascent cosmos.

The same idea applies to all other examples. The reason why tossing the newly unbound pages of *War and Peace* into the air results in a state of higher entropy is that they *began* in such a highly ordered, low entropy form. Their initial ordered form made them ripe for entropy increase. By contrast, if the pages initially were totally out of numerical order, tossing them in the air would hardly make a difference, as far as entropy goes. So the question, once again, is: how did they become so ordered? Well, Tolstoy wrote them to be presented in that order and the printer and binder followed his instructions. And the highly ordered bodies and minds of Tolstoy and the book producers, which allowed them, in turn, to create a volume of such high order, can be explained by following the same chain of reasoning we just followed for an egg, once again leading us back to the big bang. How about the partially melted ice cubes you saw at 10:30 p.m.? Now that we are trusting memories and records, you remember that just before 10 p.m. the bartender put fully formed ice cubes in your glass. He

got the ice cubes from a freezer, which was designed by a clever engineer and fabricated by talented machinists, all of whom are capable of creating something of such high order because they themselves are highly ordered life forms. And again, we can sequentially trace their order back to the highly ordered origin of the universe.

The Critical Input

The revelation we've come to is that we can trust our memories of a past with lower, not higher, entropy only if the big bang—the process, event, or happening that brought the universe into existence—started off the universe in an extraordinarily special, highly ordered state of low entropy. Without that critical input, our earlier realization that entropy should increase toward both the future and the past from any given moment would lead us to conclude that all the order we see arose from a chance fluctuation from an ordinary disordered state of high entropy, a conclusion, as we've seen, that undermines the very reasoning on which it's based. But by including the unlikely, low-entropy starting point of the universe in our analysis, we now see that the correct conclusion is that entropy increases toward the future, since probabilistic reasoning operates fully and without constraint in that direction; but entropy does not increase toward the past, since *that* use of probability would run afoul of our new proviso that the universe began in a state of low, not high, entropy.[17] Thus, conditions at the birth of the universe are critical to directing time's arrow. *The future is indeed the direction of increasing entropy. The arrow of time—the fact that things start like* this *and end like* that *but never start like* that *and end like* this*—began its flight in the highly ordered, low-entropy state of the universe at its inception.*[18]

The Remaining Puzzle

That the early universe set the direction of time's arrow is a wonderful and satisfying conclusion, but we are not done. A huge puzzle remains. How is it that the universe began in such a highly ordered configuration, setting things up so that for billions of years to follow everything could slowly evolve through steadily less ordered configurations toward higher and higher entropy? Don't lose sight of how remarkable this is. As we emphasized, from the standpoint of probability it is much more likely that the partially melted ice cubes you saw at 10:30 p.m. got there because a statistical fluke acted itself out in a glass of liquid water, than that they originated in the even less likely state of fully formed ice cubes. And what's true for ice cubes is true a gazillion times over for the whole universe. Probabilistically speaking, it is mind-bogglingly more likely that everything we now see in the universe arose from a rare but every-so-often-expectable statistical aberration away from total disorder, rather than having slowly evolved from the even more unlikely, the incredibly more ordered, the astoundingly low-entropy starting point required by the big bang.[19]

Yet, when we went with the odds and imagined that everything popped into existence by a statistical fluke, we found ourselves in a quagmire: that route called into question the laws of physics themselves. And so we are inclined to buck the bookies and go with a low-entropy big bang as the explanation for the arrow of time. The puzzle then is to explain how the universe began in such an unlikely, highly ordered configuration. *That* is the question to which the arrow of time points. It all comes down to cosmology.[20]

We will take up a detailed discussion of cosmology in Chapters 8 through 11, but notice first that our discussion of time suffers from a serious shortcoming: everything we've said has been based purely on classical physics. Let's now consider how quantum mechanics affects our understanding of time and our pursuit of its arrow.