

What Is Life?

Erwin Schrödinger (1887–1961), reluctant cofounder of quantum mechanics, 1933 Nobel laureate in physics, and author of a famous thought experiment involving cruelty to felines, was used to speaking his mind. So much so that, after the Nazis came to power in 1933, he resigned from his chair at the University of Berlin, which he had taken over from Max Planck (1858–1947; 1918 Nobel laureate in physics) just six years earlier, and emigrated first to Oxford, then to his native Austria, from where he was exiled again after the *Anschluss*. In 1939, the government of neutral Ireland invited him to take up a chair of theoretical physics at the newly founded Dublin Institute for Advanced Studies. Although he was a political refugee, his landing was a soft one and he greatly benefited from his time at Dublin, where he remained for the next seventeen years.

One of the obligations of Schrödinger's new job was an annual public lecture. In 1943, he held a series of three lectures at Trinity College Dublin, where an audience of more than four hundred heard him discourse on the topic "What Is Life?" At a time when there was no such thing as biophysics, this venture of a theoretical physicist into the domain of biologists was unprecedented. Moreover, there was virtually nothing known in biology that would have satisfied the strict thinking of a physicist. So instead of giving answers, Schrödinger formulated some fundamental questions of biology, as seen by a physicist.

Schrödinger mainly covered two fundamental aspects of life, namely heredity and thermodynamics. He framed these in the basic questions of how life creates "order from order" and how it creates "order from disorder." In his analysis of genetics (order from order), he estimated the number of atoms contained in a gene (then a highly abstract concept). He proposed that the genetic information might be encoded in something resembling an aperiodic crystal—that is, a combination of a regular structure with information-bearing variations—an idea that, in retrospect, seems startlingly prescient. In the second half of his discourse, Schrödinger clarified that organisms can create ordered arrangements of molecules (and cells and tissues, in the case of higher organisms) within themselves, by creating even greater disorder in the environment. Thus

was the evolution of highly complex organisms from a chaotic pool of simple, lifeless chemicals kept in line with the second law of thermodynamics.

Ultimately Schrödinger's lectures were published as a small book—which is still in print today—that was hugely influential. For the first time a prominent scientist had raised the question of how the *physics* of our Universe fundamentally constrains its *biology*. Still, in 1943 the question “what is life?” was wide open and posed major challenges not just to biology but across all of science. In the six decades that have passed since then, many aspects of this question have been resolved, such that today, in this opening chapter, we can take a stab not only at defining what life is but also at listing some of its most fundamental requirements.

Life

So, with the knowledge we have at the beginning of the twenty-first century, what, precisely, is life? Most answers to this question are reminiscent of the claim of U.S. Supreme Court Justice Potter Stewart (1915–85) that, while he could not precisely and unambiguously define pornography, “I know it when I see it.” But that kind of empirical approach, of course, does not get us very far if we are going to embark on a deep and rigorous evaluation of the origins of life and its relationship to the Universe at large.

Moreover, if we are interested in what *might* have happened—the range of possibilities that could have unfolded on Earth, or might be happening elsewhere in the Universe—we have to attempt to define the boundary conditions of life. That is, we need to attempt to understand the range of conditions and events that had to conspire to make life possible.* In this, as perhaps in the definition of life itself, we must necessarily be somewhat parochial; our understanding of the conditions under which life can arise and evolve is almost certainly going to be flavored by deeply held preconceptions based on our understanding of life on Earth. But as long as we are aware of this underlying bias, we can at least tackle each of the seemingly necessary conditions in as unbiased and logical a fashion as is (even the word itself is telling) *humanly* possible.

Definition of Life

Life scientists should know what the first word in their job title means, but practitioners of various disciplines ranging from the origins of life

*Keep in mind, too, that evolution is quite good at generating organisms that are fit to survive under very harsh conditions. And thus the set of conditions under which life can thrive is probably much broader than the set of conditions under which it can arise from inanimate matter in the first place.

through to modern astrobiology have consistently failed to come up with an all-inclusive definition of life. Nevertheless, we shall conjure up a working definition of life that, if not perfect, will suffice as a basis for our discussion.

The most striking property that distinguishes living systems from the inanimate world is their ability to copy themselves, a process scientifically described by the term *self-replicating*. Among *Homo sapiens* the process is more colloquially captured in the phrase "get married, settle down, and have kids." The fact that living things copy themselves is so central to all of biology that some wag once pointed out that "life is just a DNA molecule's way of reproducing itself." All of biology, from bacterial mats through to warring nations, can be described as tools for or consequences of the replication of genes.

Another key limit to our discussion is to define life as a *chemical system* (as opposed to a mechanical or electronic system). Over decades, writers of science fiction and of putative nonfiction extrapolating current (nano)technological trends into the future have suggested that self-replicating, microscopically small robots will soon be cleaning out our arteries, degrading toxic waste, and generally making themselves useful. Irrespective of the accuracy of these predictions, it seems likely that physical laws of the Universe allow the creation of mechanical beings that can construct copies of themselves and thus meet our first criterion for life. Similarly, there are viable organisms in cyberspace, known as worms. While they require a computer, an internet connection, and, typically, some poorly written software in order to reproduce, one might argue that these items constitute their ecosystem; we are just as critically reliant on our ecosystem for reproduction. When working out a universal definition of life, it's not so easy to dismiss these potentially living things out of hand. Especially when, as technology progresses, the boundaries between biological, mechanical, and electronic systems will probably slowly erode, as brains will be interfaced with computers and micro-robots will resemble insects.

Considering the origins and distribution of life in the Universe, however, it is difficult to imagine that mechanoid life (much less life dependent on the existence of an internet) would have arisen spontaneously. The problem is that mechanical things, by definition, use parts that are larger than molecules (if a system consists of molecular-scale parts, then it is by definition a *chemical system*). Before the creation of the first organisms, these parts would have to be moved around by the random fluctuations of solvent molecules moving to and fro, which is called Brownian motion. Brownian motion isn't all that fast: if you gently open a bottle of perfume, how long does it take Brownian motion to waft the

molecules of fragrance across the room? And since thermal motions vary with the square root of the mass of the diffusing object, it would take far longer than the age of the Universe for a bucket of watch parts to spontaneously assemble into a watch, much less into a machine capable of copying itself. Thus, while mechanically based life might arise by the intelligent design of chemical life forms, it seems unlikely that it can arise spontaneously. In a nutshell, it seems fair not to worry about whether the self-replicating robots of the planet Lexus Nine are alive.

A final, but critical, element in our definition of life emerges from the observation that not all self-replicating chemical systems are alive. For example, crystals are, in a sense, self-replicating. This is particularly true in a supersaturated solution. Under such conditions, if one were to smash a growing crystal into smaller pieces, each of the pieces would in turn grow into a new and larger crystal. Crystals even breed true. For example, whereas individual molecules of sodium chlorate are not "handed" (they are superimposable on their mirror images—more on this in chapter 5), crystals of this substance do have a handedness (chirality). When you allow a sodium chlorate solution to crystallize, half of the crystals will be the mirror image of the other half: half left handed, the other half right. However, if you take a supersaturated solution of sodium chlorate and stir vigorously while it begins to crystallize, all of the crystals that form will be either right or left handed. Why is this? It is because the vigorous stirring shatters the first crystal that forms, and the minicrystals thus formed nucleate the growth of all the crystals that follow, causing them to adopt the same handedness. Crystallization is self-replication. Equally clearly, though, while "a diamond is forever," it is not and, with rare exceptions,* never was alive. So what do we need to add to our definition of life in order to discriminate between inanimate crystals and truly animate chemical systems? In a word, evolution.

Living beings produce offspring in their own image by the replication of their genetic material. But the replication is not perfect: random genetic mutations produce inheritable differences that may improve or impede an offspring's viability. These give natural selection a chance to shape the fate of future generations and, indeed, the evolution of a species. Evolution, with the inherent adaptation under selective pressures, is a fundamental property of life and clearly distinguishes it from inanimate, if sometimes self-replicating, materials. A crystal makes per-

*"Rare exceptions," you dare ask? In the 1950s, shortly after General Electric developed a commercially viable method for the production of industrial diamonds, the technicians there synthesized some out of peanut butter. Upping the ante, the company LifeGems now provides the service of converting some of the carbon in your deceased "loved ones" into permanent and rather sparkly keepsakes.

fect copies of itself. The first crystal of quartz that condensed out of the solar nebula 4.57 billion years ago is identical to the quartz that crystallized last week in the Corning Glassware plant in upstate New York. Crystals and crystallization are changeless, incapable of evolving into new, more complex, and better forms. And thus they are not, and never have been, alive.

Limitations of Our Definition

So, there we have our definition. Life is a self-replicating chemical system capable of evolving such that its offspring might be better suited for survival. As definitions go, this one is nice, clean and concise. Too bad it is fatally flawed. Or at least seriously limited. While a chemical system that is capable of reproduction and evolution is clearly alive, the reverse is not necessarily true; many things that fail to meet these criteria are obviously *also* alive. Those of us whose child-bearing years are past, for example, might take umbrage at the suggestion that they are not alive simply because they are no longer reproducing. But while it would be nice to have a definition of life that could be applied as a litmus test to every single specimen and include even post-reproductive academics, it is not necessary for our discussion. In many regards, evolution acts on the level of populations and species. For a species to thrive it must have individuals capable of reproducing, but there may very well be members of the species that serve its survival without reproducing at all, as is true for most members of ant or bee colonies. Thus, by defining a living organism as a self-replicating, evolving chemical system, we cover all species known to date (if not all individuals). Moreover, since replicating organisms must have preceded any given nonreplicating organism, this definition is sufficient for our needs because it does not artificially constrain our discussion of the origins and evolution of life.

Requirements for Life

What, then, are the fundamental conditions that life requires? Given that our knowledge of this subject is necessarily parochial, we should cast our net wide, making an effort not to mistake Terrestrial constraints for universal ones. Still, there are a number of criteria that seem to be absolutely critical elements for the formation of life.

Life requires chemistry. This means that life requires atoms more complex than hydrogen, whose solo chemistry is limited to the formation of H_2 , and helium, which is one of the few chemical elements that lack any chemistry whatsoever. Even taking into account our potentially parochial, Terrestrial biases, it seems fairly certain that a self-replicating chemical system cannot be built using just the reaction $H + H \rightarrow H_2$.

*	Ce	Pr	Nd	Pm	Sm	Eu	Gd	Tb	Dy	Ho	Er	Tm	Yb	Lu
**	Th	Pa	U	Np	Pu	Am	Cm	Bk	Cf	Es	Fm	Md	No	Lr

Thus the formation and evolution of chemical life will require atoms more complex than the two lightest atoms. And, as we will see, while hydrogen and helium were formed in great abundance in the first minutes of the Universe, the formation of heavier atoms was a far more delicate matter.

It seems a fair assumption that a chemical system capable of copying itself will require at least a modest degree of complexity, and building complex molecules requires that we bond many atoms together. Clearly this cannot be done for the noble gases helium (He), neon (Ne), and argon (Ar), as these atoms do not participate in any chemistry. Nor

can we build a complex chemistry based on atoms, such as chlorine, that make only one bond; at best they can form diatomics such as the aforementioned H_2 . Thus in our search for atoms that could serve as the framework chemistry of life we can discount the first, second-to-last, and last columns of the periodic table, which are filled with such "uninteresting" atoms.

Similarly, to serve as the foundation of complex molecules, an atom must form very strong bonds to other atoms, and probably to itself (more precisely, to another atom of the same type). What do we mean by strong? We mean bonds that are hundreds of times stronger than the energy contained in a typical molecular collision, lest these same collisions tear the molecules apart. As we go down the periodic table, the outer electrons in each succeeding row of atoms—the electrons that participate in bonds—are more and more weakly bound to the nucleus. This occurs because each succeeding row in the table represents another filled shell of electrons, and with each row the outer electrons are more and more shielded from (i.e., less and less attracted to) the positively charged nucleus. Because of this, the bonding strength of the second, third, and fourth rows of the periodic table becomes progressively weaker. This is a serious issue. Whereas carbon, boron, and the like make for long, extremely strong chains of molecules (e.g., the long polymer chains that plastics are made of), no one has ever made a chain of silicon that was more than two atoms long; the Si-Si bond is simply too weak. Only the second-row elements are capable of forming strong covalent bonds to one another and to elements in the other rows. Thus we are probably on fairly strong ground in discounting all but this second row of elements in our quest to find the minimum set of materials necessary to form life.

There is one last criterion that might segregate reasonable life-forming elements from those that are much less likely to participate in the process: abundance. Even a casual glance around the Earth suggests that some elements are much more precious than others; gold is expensive because it is rare, whereas oxygen costs just a few cents per kilogram in industrial quantities. We go into this in detail in chapter 3, so let it suffice to say here that among the eight elements in the second row of the periodic table, lithium, beryllium, boron, and fluorine are relatively rare in the Universe (fig. 1.2). Thus a theory about alternative life forms that relies critically on these elements is significantly more suspicious than one that does not.

Life based on molecules almost certainly requires a *solvent* in which to move them around. Because mass transport through solids is at best extremely slow, solid-phase chemistry is far too limited to support the




FIG. 1.2. The relative abundances of elements in the Sun map out the relative abundances available in the presolar nebula and thus, in turn, available to serve as the basic building blocks of life. The relative abundances of elements such as carbon and oxygen render them inherently more likely to serve as the basis of life than the rarer elements. Of note, the Sun is significantly enriched in elements heavier than helium relative to other stars its size. Nevertheless, this qualitative pattern is common among stars of the Sun's generation throughout the Universe (more on this in the next chapter).

complex networks of chemistry required for a self-replicating organism. This observation once again highlights the unique ability of the second-row elements to support life; Terrestrial animals eat water-soluble carbon compounds, such as the sugar and other carbohydrates in your morning doughnut, oxidize them with water-soluble oxygen, and excrete equally water-soluble carbon dioxide; silicon-based life forms, in contrast, would have a much more difficult time exhaling silicon dioxide, which tends to appear in solid forms, such as sand.

The human body contains around 70% water, highlighting the fact that, for Terrestrial organisms, the solvent in question is water. But is water the only plausible solvent for life? Once again a quick glance at the periodic table suggests that out of the (very finite) list of potential "biotic solvents" water may well be the only reasonable option. Water has

so many properties that render it ideally suited as a biological solvent that its ability to form the basis of biochemistry may well be unique.

Some of the “ideal” properties of water are well known, and others, while less so, are no less critical for life on Earth. An example is taught to almost every elementary school student: water is one of the very few substances that expand when they freeze. Because of this, ice floats. If, instead, the ocean were filled with liquid ammonia, its winter pack “ice” would sink, where it would be insulated from the summer’s warmth and prevented from seasonally melting. With each passing year, more and more of the ocean’s volume would be locked up in the solid until, in a timeframe quite rapid by geological standards, the planet would freeze over, with only a thin seasonal layer of liquid on the surface. Could such a frozen ocean support the origins of life? Perhaps. But a permanently liquid ocean, with its ability to transport nutrients and modulate temperature, seems more likely to do the trick.

Water also has an extraordinary ability to absorb heat without much of a rise in its temperature, which is why we use it as a carrier for heat in central heating systems and hot water bottles and, conversely, also as a coolant. In more precise terms, the heat capacity of liquid water is 1 cal/g °C (by definition, it takes precisely 1 calorie—that is, 4.184 joules—to heat 1 gram of water from 14.5°C to 15.5°C). This value is about three times higher than that of typical rock or metal. This high heat capacity helps to moderate the Earth’s climate, a seemingly critical event in the origins and evolution of life that we will cover in more detail in chapter 3. On a related note, thanks to its unique ability to form extended hydrogen-bonding networks, water remains liquid over a surprisingly broad, hundred-degree Celsius temperature range, thus helping to ensure that, even if the climate does fluctuate radically, liquid solvent will be available for life.

In addition to these important physical properties, the chemical properties of water seem to render it ideally suited as the basis for life. For example, the dielectric constant of water is around 80, which is significantly higher than that of any other cosmologically abundant liquid. This means that two oppositely charged ions in water are attracted with one-eightieth the force they would feel in a vacuum. Because of this, water can shield charged ions from one another, allowing them to be readily taken into solution, where they can perform chemistry. Water also has the highest molar density of any molecular liquid; fully 55.5 moles of water (3.3×10^{25} molecules) are crammed in each and every liter of the stuff. No other liquid packs anywhere near this many molecules in a given volume. Because of this extraordinary molar density, the entropy

cost of "organizing" water ("the solvent") around any molecules dissolved in it ("the solute") is quite high (many water molecules need to be moved out of the way to make room for each cubic nanometer of solute), and thus water tends to force many types of dissolved molecules to organize themselves in order to minimize this entropic cost. This organizing effect, which is called the "hydrophobic effect,"* plays a critical role in organizing biomolecules on Earth. Lastly, water is cosmologically abundant, as its components, hydrogen and oxygen, are the first and third most abundant atoms in the Universe (fig. 1.2).

Of course, the fact that water is well suited for life on Earth doesn't automatically rule out that life elsewhere might be based on some other solvent. Or does it? It would be hard to find an alternative, as no other liquid has even a fraction of the favorable attributes of water. Hydrogen fluoride perhaps comes closest. Compared with water, it has a slightly higher dielectric constant (84, to water's 80), and thus it is at least as good a solvent for ionic materials. It also has a slightly wider, 102 degrees Celsius, liquid range (at atmospheric pressure it freezes at -83°C and boils at 19°C) and a comparable molar density (48.0 versus 55.5 mol/L). But as fluorine is cosmologically rare—it is about 1/100,000 as abundant as oxygen—it seems very unlikely that there are little purple fish happily swimming in seas of liquid HF on the planet Zap Seven.

Thus we are probably safe ruling out hydrogen fluoride. And none of the other molecular liquids formed by the cosmologically abundant elements (such as ammonia, hydrogen sulfide, or methane) comes anywhere near as close to the ideal properties of water as does HF; their liquid ranges are extremely small, their ability to solvate ionic materials is poor to effectively nonexistent, and their ability to regulate climate is extremely limited (table 1.1). From such considerations emerges the near certainty not only that life has an absolute requirement for a liquid solvent, but that water is by far the most "qualified" solvent to fulfill that role. This is not to say that life cannot have arisen based on other solvents; simply that the origins of life face a much more significant hurdle in the absence of this remarkable and abundant liquid.

Life also requires a solid or liquid *substrate*. The reason that life probably cannot exist in the gas phase is that molecules of sufficient complexity to form life are inevitably too dense to stay suspended. This is, of course, a much bigger constraint on the origins of life than on its ability to thrive after it has arisen; if the surface of the Earth were slowly to become uninhabitable, it is a pretty good bet that at least some bacteria would adapt to full-time living in cloud droplets. Indeed, as we discuss

*It's why, for example, "oil and water don't mix."

TABLE 1.1
Physical properties of potential biological solvents

Solvent	Formula	Liquid range (°C, at 1 atmosphere)	Molar density (mol/L)	Heat capacity (cal/g °C)	Dielectric constant
Water	H ₂ O	0 to +100	55.5	1.0	80
Hydrogen fluoride	HF	-83 to +19	48.0	0.8	84
Ammonia	NH ₃	-78 to -34	40.0	1.1	25
Hydrogen sulfide	H ₂ S	-85 to -6	26.8	0.5	9
Methane	CH ₄	-182 to -161	26.4	0.7	2
Hydrogen	H ₂	-259 to -253	35.0	0.002	1

in a later chapter, some may already have done so. But the limited mass transport and limited size of condensed bodies that can occur in the gas phase make this realm an exceedingly unlikely one for the *origins* of life. This effectively rules out life on Jupiter (if life couldn't get started in the first place, it was unlikely to have evolved into giant, hydrogen-filled Hindenburgoids), much less, for example, in interstellar space.*

And let us not forget that life requires *energy*. This is obvious for the chemist, as living organisms create an implausible amount of order out of disorder, such as when the randomly distributed molecules of carbon dioxide and fertilizers end up in the highly nonrandom structure of a plant. According to the second law of thermodynamics, they can achieve this only if at least as much entropy (roughly speaking, a measure of molecular disorder) is created elsewhere. By using energy from an external source, life can swap energy for entropy: the living organisms get the calories and the order, while the rest of the Universe pays the price.

Thus, life requires an external disequilibrium (an "ordered" state) whose tendency to drive chemical reactions toward a more equilibrated ("disordered") state it can exploit for its own purpose of organizing its molecules into some pattern capable of reproduction. Among the most abundant sources of disequilibrium in the Universe are temperature differences: the fact that stars are much hotter than the Universe at large. Because of this, the copious number of high-energy photons emitted by a star can be absorbed by the surface of a (much cooler) planet. Here on

*This didn't stop the late cosmologist Fred Hoyle from writing a wonderful bit of science fiction, *The Black Cloud*, about an intelligent interstellar cloud that pays the Solar System a visit and accidentally wipes out most of humanity. It simply hadn't occurred to the cloud that life was possible on the cold, densely packed surface of a planet.

Earth, plants take advantage of this disequilibrium and use it to feed the striking disequilibrium that is our biosphere. For example, the presence of combustible wood in an atmosphere containing oxygen is a clear deviation from chemical equilibrium with respect to a mixture of water and carbon dioxide, as forest fires remind us. We animals, in turn, take advantage of the latter disequilibrium when we oxidize the carbohydrates in our morning doughnut to generate the energy we use to run our metabolic processes.

Substrates and solvents and thermodynamics aside, life also presumably requires *time*. And the narrower the range of conditions under which life can arise in the first place, the more unlikely will be the occurrence of a sufficiently stable environment that will stay within the range for sufficiently long. The Universe is a dangerous place. The luminosity of a star changes, and with it the temperature of any planets warmed by its light. Planets are sometimes struck by asteroids so large that the energy imparted by the impact can boil oceans and sterilize worlds. Atmospheres escape into space. Rotational axes tilt, plunging planets into million-year winters. Supernovae explode with the power of a billion suns, sterilizing any planet within a few hundred light-years. Considering these risks, it is clear that not all of the environments in the Universe that are capable of supporting the formation of life will remain stable long enough for life to arise at all, much less gain a secure footing.

Conclusions

So the recipe for life to arise somewhere in the Universe seems relatively straightforward. All we need is some water, carbon on a solid (or liquid) planetary surface, an energy source, some time, and we're off. But is it that easy? What is required to produce a water-and-carbon-bearing planetary environment that provides energy sources and yet is stable over eons? And how often are these conditions met? And if we find these conditions, how likely is it that life will arise? The following chapters explore each of these critical questions in turn.

More than fifty years after Schrödinger's lectures, the most fundamental aspects of the questions he asked have been answered, even if some details are left to fill in. In the summer of 1993, a dozen prominent scientists, including Nobel laureates Christian de Duve and Manfred Eigen, science popularizers Jared Diamond and Stephen Jay Gould (1941–2002), and evolutionary pioneers John Maynard Smith and Leslie Orgel, assembled at Trinity College Dublin to commemorate the lectures and to deliver new lectures. Their ambitious goal was to set a research agenda for the next fifty years of life science research, as Schrödinger had done. Because the investigation of present life on our planet has become

relatively straightforward, many of the lectures focused on the mysteries of the origins and early evolution of life on Earth, which will also loom large in the chapters around the middle of this book.

Further Reading

Expectation of and constraints on life in the Universe. Schulze-Makuch, Dirk, and Irwin, Louis. *Life in the Universe*. Berlin: Springer-Verlag, 2004.

The history of origins-of-life research. Fry, I. *The Emergence of Life on Earth*. New Brunswick, NJ: Rutgers University Press, 2000.