

A reprint from
American Scientist
the magazine of Sigma Xi, The Scientific Research Society

This reprint is provided for personal and noncommercial use. For any other use, please send a request to Permissions, American Scientist, P.O. Box 13975, Research Triangle Park, NC, 27709, U.S.A., or by electronic mail to perms@amsci.org. ©Sigma Xi, The Scientific Research Society and other rightsholders

The Origin of Life

A case is made for the descent of electrons

James Trefil, Harold J. Morowitz and Eric Smith

As the frontiers of knowledge have advanced, scientists have resolved one creation question after another. We now have a pretty good understanding of the origin of the Sun and the Earth, and cosmologists can take us to within a fraction of a second of the beginning of the universe itself. We know how life, once it began, was able to proliferate and diversify until it filled (and in many cases created) every niche on the planet. Yet one of the most obvious big questions—how did life arise from inorganic matter?—remains a great unknown.

Our progress on this question has been impeded by a formidable cognitive barrier. Because we perceive a deep gap when we think about the difference between inorganic matter and life, we feel that nature must have made a big leap to cross that gap. This point of view has led to searches for ways large and complex molecules could have formed early in Earth's history, a daunting task. The essential problem is that in modern living systems, chemical reactions in cells are mediated by protein catalysts called enzymes. The information encoded in the nucleic acids DNA and RNA is required to make the proteins; yet the proteins are required to make the nucleic acids. Furthermore, both proteins and nucleic acids are large molecules consisting of strings of small component molecules whose synthesis is supervised by proteins and nucleic acids. We have two chickens, two eggs,

and no answer to the old problem of which came first.

In this article we present a view gaining attention in the origin-of-life community that takes the question out of the hatchery and places it squarely in the realm of accessible, plausible chemistry. As we see it, the early steps on the way to life are an inevitable, incremental result of the operation of the laws of chemistry and physics operating under the conditions that existed on the early Earth, a result that can be understood in terms of known (or at least knowable) laws of nature. As such, the early stages in the emergence of life are no more surprising, no more accidental, than water flowing downhill.

The new approach requires that we adopt new ways of looking at two important fields of science. As we will see below, we will have to adjust our view of both cellular biochemistry and thermodynamics. Before we talk about these new ideas, however, it will be useful to place them in context by outlining a little of the history of research on the origin of life.

The Origin of Origins

Most historians would say that the modern era of experimental research in origin-of-life studies began in a basement laboratory in the chemistry department of the University of Chicago in 1953. Harold Urey, a Nobel laureate in chemistry, and Stanley Miller, then a graduate student, put together a tabletop apparatus designed to look at the kinds of chemical processes that might have occurred on the planet soon after its birth. They showed that organic molecules (in this case amino acids) could be created from inorganic materials by natural environmental conditions such as acidic solution, heat and electrical discharge

(lightning), without the mediation of enzymes. This finding triggered a wave of new thinking about both the origin and nature of life. (Today, the consensus is that Miller and Urey had the wrong atmospheric components in their apparatus, so the process they discovered was probably not representative of the emergence of life on Earth. It nevertheless pointed to the potential fecundity and diversity of nonenzymatic primordial chemistry.)

Since 1953, we have found many of the same simple organic molecules in meteorites, comets and even interstellar gas clouds. Far from being special, then, the simplest of the molecules we find in living systems—life's building blocks—seem to be quite common in nature. To many, the real question was how these basic building blocks got put together into living systems, and, equally important, how the molecules that led to modern life were selected out of the messy molecular milieu in which they arose.

The ubiquity of simple molecules suggested an appealing scenario that had a profound effect on the way investigators approached the origin of life throughout the last half of the 20th century. The scenario went like this: After the Earth cooled enough to allow oceans to form, the Miller-Urey process or something like it produced a rain of organic matter. In a relatively short time, the ocean became a broth of these molecules, and given enough time, the right combination of molecules came together by pure chance to form a replicating entity of some kind that evolved into modern life.

Scientists called this scenario the Oparin-Haldane conjecture, but it was given a provocative nickname that endures in the popular consciousness—Primordial Soup.

James Trefil is Clarence J. Robinson Professor of Physics at George Mason University. Harold Morowitz is Clarence J. Robinson Professor of Biology and Natural Philosophy at George Mason University. Eric Smith is a professor at the Santa Fe Institute. Address for Trefil: 207F East Building, MS 1D6, George Mason University, Fairfax, VA 22030. Internet: jtrefil@gmu.edu.



Figure 1. Origin-of-life studies became an experimental science with the Miller-Urey experiment, which produced organic molecules in a flask from components thought to be present in the paleogeological atmosphere—homemade primordial soup. But how did soup ingredients become life? A recent model called Metabolism First proposes that life didn't climb over a thermodynamic barrier, it fell into place according to knowable laws of chemistry and thermodynamics. (Image courtesy of Scripps Institution of Oceanography, University of California, San Diego.)

The essential legacy of the Primordial Soup was twofold: It simplified the notion of the origin of life to a single pivotal event, and then it proposed that that event—the step that occurred after the molecules were made—was a result of chance. In the standard language, life is to be seen, in the end, as a “frozen accident.” In this view, many fundamental details about the structure of life are not amenable to explanation. The architecture of life is just one of those things. Although many modern theories are less extreme than this, frozen-accident thinking still influences what some of us ask about the origin of life and how we prioritize our experiments.

RNA World

The next major advance came in the early 1980s, when Thomas Cech and Sidney Altman showed that some RNA molecules can act as enzyme-like

catalysts. The frozen-accident argument was then replaced by a suggestive scenario in which something like RNA was assembled by chance, and was then able to fill twin roles as both enzyme and hereditary molecule in the runup to life. The RNA systems were then acted upon by natural selection, leading to greater molecular complexity and, eventually, something like modern life. Whereas most scientists believe, on the basis of Cech and Altman's work, that life went through an early RNA-dominated phase (dubbed “RNA World”), the “RNA First” scenario has again a quality of frozen accident. Between prebiological chemistry and RNA World, a large leap occurs, requiring that molecules appear having at least a familial resemblance to the complex molecules in the vials of Cech and Altman, for that is the assumption upon which

the relevance of their findings to the origin of life depends.

Inserting RNA molecules into an RNA First scenario without explaining how they got there seems to us an inadequate foundation for an origin theory. The RNA molecule is too complex, requiring assembly first of the monomeric constituents of RNA, then assembly of strings of monomers into polymers. As a random event without a highly structured chemical context, this sequence has a forbiddingly low probability and the process lacks a plausible chemical explanation, despite considerable effort to supply one. We find it more natural to infer that by the time complex RNA was possible, life was already well on the road to complexity. We believe further that we can see the primordial chemical architecture preserved in the universal metabolic chemistry we observe today.

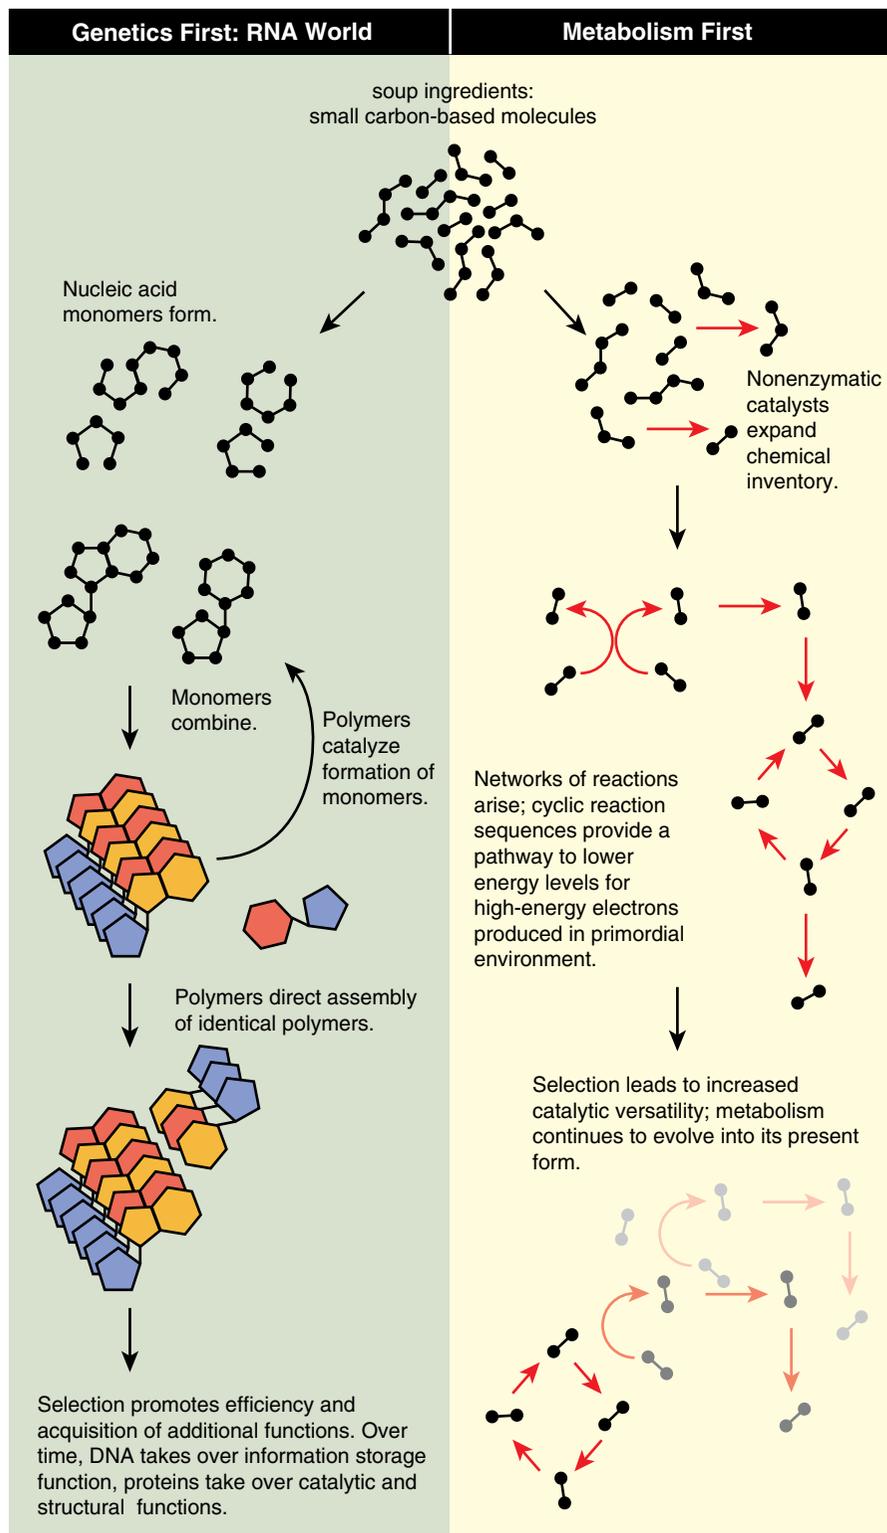


Figure 2. RNA World has been the prevailing theory for the origin of life since the 1980s. The emergence of a self-replicating catalytic molecule accounts for signature capabilities of living systems, but it doesn't explain how the protobiological molecule itself arose. Metabolism First seeks the answer in primitive reaction networks that generate their own constituents, offering a substrate for chemical selection and a launchpad for life.

The compelling feature of RNA World is that a primordial molecule provided both catalytic power and the ability to propagate its chemical identity over generations. As the catalytic versatil-

ity of these primordial RNA molecules increased due to random variation and selection, metabolic complexity began to emerge. From that stage, RNA had roles in both control of metabolism and

continuity across generations, as it does today. Depending on which function one prefers to emphasize, these overall models have been called "Control First" or "Genetics First." In either case, the proliferation of metabolism depended on RNA being there first.

Adherents have come to call the other possibility "Metabolism First" (though by this they have meant many slightly different things). In our version of Metabolism First, the earliest steps toward life required neither DNA nor RNA, and may not even have involved spatial compartments like cells; the earliest reactions could have occurred in the voids of porous rock, perhaps filled with organic gels deposited as suggested in the Oparin-Haldane model. We believe this early version of metabolism consisted of a series of simple chemical reactions running without the aid of complex enzymes, via the catalytic action of networks of small molecules, perhaps aided by naturally occurring minerals. If the network generated its own constituents—if it was recursive—it could serve as the core of a self-amplifying chemical system subject to selection. We propose that such a system arose and that much of that early core remains as the universal part of modern biochemistry, the reaction sequences shared by all living beings. Further elaborations would have been added to it as cells formed and came under RNA control, and as organisms specialized as participants in more complex ecosystems.

Networks of synthetic pathways that are recursive and self-catalyzing are widely known in organic chemistry, but they are notorious for generating a mass of side products, which may disrupt the reaction system or simply dilute the reactants, preventing them from accumulating within a pathway. The important feature necessary for chemical selection in such a network, which remains to be demonstrated, is feedback-driven self-pruning of side reactions, resulting in a limited suite of pathways capable of concentrating reagents as metabolism does. The search for such self-pruning is one of the most actively pursued research fronts in Metabolism First research.

A Pair of Analogies

Here's an analogy that will provide an outline for the argument we make: Consider the requirements of the U.S. Interstate highway system. The system

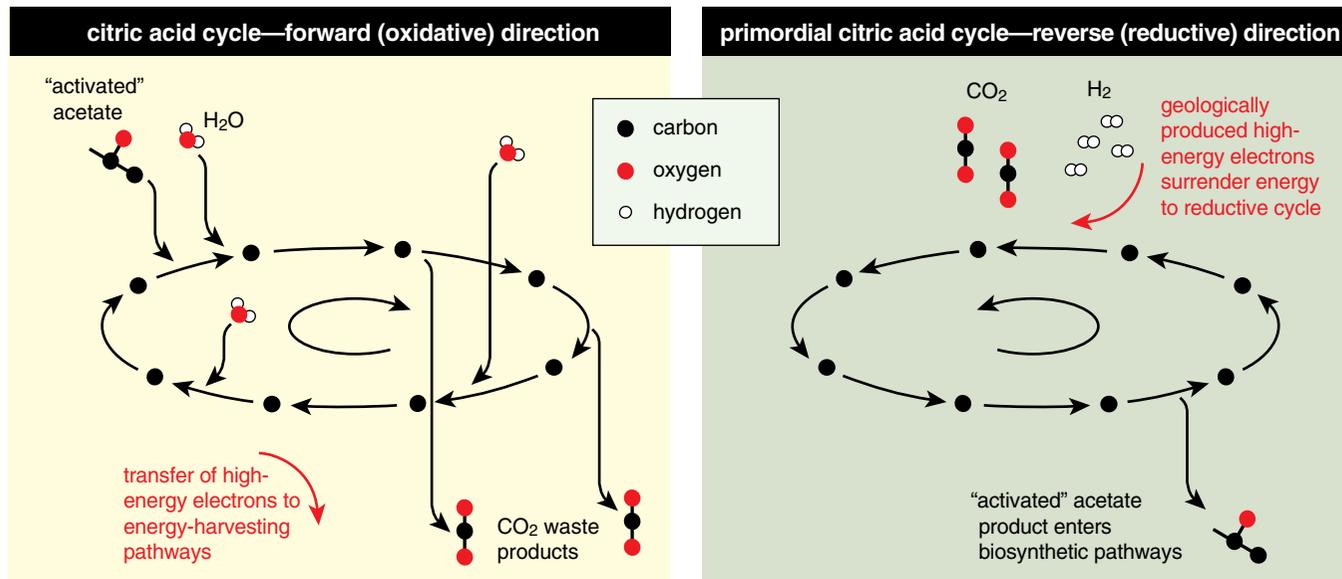


Figure 3. The reactions and molecules of the citric acid cycle are universal in modern organisms. However, in many microbial species, the cycle runs in reverse. Instead of oxidizing the fuel molecule acetate (“activated” by attachment to a carrier molecule) and releasing CO₂ as waste, the reverse cycle incorporates CO₂ in organic molecules by exploiting the electron-transfer potential (“reducing power”) of geologically produced molecules such as H₂. A reductive cycle could have served as the foundation for primordial biosynthesis.

Imagine a large pond of water sitting on top of a hill. We know that there are any number of other states—any in which the water is lower than it is at the top—which have lower energy and are therefore states toward which the system will tend to evolve over time. In terms of our question, the “problem” faced by the system is how to get water from its initial state to any state of lower energy—how to get the water down the hill. We need not think of the laws of physics as being endpoint directed; rather, they simply adjudicate between states of higher or lower energy, with a preference for lower. Can we apply the same reasoning to the chemistry of life?

For real hills, we understand not only that the water will flow downward but also many things about how it will do so. Molecules of water will not each flow down a random path. Instead the flowing water will cut a channel in the hillside. In fact, the flow of water is at once constructing a channel and contributing to the collapse of the energy imbalance that drives the entire process. In addition, if we look at this process in detail, we see that what really matters is the configuration of the earth near the top of the hill, for it is there that the channeling process starts. This part of the analogy turns out to be particularly appropriate when we consider early chemical reactions.

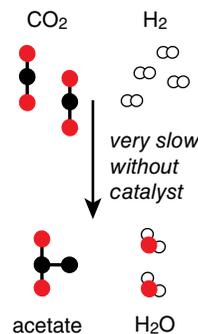
In the analogy, the “problem” is the fact that the water begins in a state of high energy; the creation of the chan-

nel “solves” this problem by allowing the water to move to a lower energy state. Furthermore, the dynamics of the system are such that once the channel is established, subsequent flow will reinforce and strengthen it. There are many such systems of channels in nature—the lightning bolt is an example, although in that case the forces at work are electrical, not gravitational. (When lightning occurs, positive and negative charges become separated between clouds and the ground. The charge separation ionizes atoms in the air, creating a conducting channel through which the charges flow—the lightning bolt—much as water flows down a hill).

We argue that the appearance of life on our planet followed the creation of just such a channel, except that it was a channel in a chemical rather than a geological landscape. In the abiotic world of the early Earth, likely in a chemically excited environment, reservoirs of energy accumulated. In effect, electrons (along with certain key ions) were pumped up chemical hills. Like the water in our analogy, those electrons possessed stored energy. The “problem” was how to release it. In the words of Albert Szent-Gyorgi: “Life is nothing but an electron looking for a place to rest.”

For example, carbon dioxide and hydrogen molecules are produced copiously in ordinary geochemical environments such as deep sea vents, creating a situation analogous to the water on the hill. The energy of this system can be

lowered if the electrons in the hydrogen “roll down the hill” by combining with the atoms of carbon dioxide in a chemical reaction that produces water and



acetate (a molecule with two carbon atoms). In the abiotic world, however, this particular reaction takes place so slowly that the electrons in the hydrogen molecules find themselves effectively stranded at the top of the energy hill.

In this example, the problem that is solved by the presence of life is getting energized electrons back down the chemical hill. This is accomplished by the establishment of a sequence of biochemical channels, each contributing to the whole. (Think of the water cutting multiple channels in the hill). The reactions that create those channels would involve simple chemical transactions between small organic molecules.

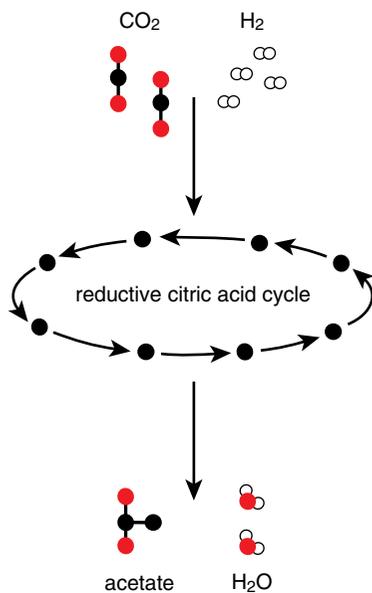
How can we translate these sorts of general arguments into a reasonable scenario for the appearance of the first living thing? One way would be to look

closely at the metabolic chart shown earlier, the diagram that maps the basic chemical reactions in all living systems.

At the very core of metabolism—the starting point for the synthetic pathways of all biomolecules—is a relatively simple set of reactions known as the citric acid cycle (also called the tricarboxylic acid cycle or the Krebs cycle). The cycle involves eight molecules, each a carboxylic acid (a molecule containing —COO groups). In most present-day life forms on Earth, the citric acid cycle operates to break organic molecules down into carbon dioxide and water, using oxygen to produce energy for the cell—in effect, “burning” those molecules as fuel. (Technically, a molecule like glucose is first broken down into smaller molecules like pyruvate, which is then fed into the citric acid cycle. Full decomposition of pyruvate to CO₂ and water is facilitated by transfer of high-energy electrons to certain coreactants that, in the modern cell, ferry the electrons to other reactions). When the cycle operates in this way, we say that it is in its oxidative mode.

The cycle can also operate in the opposite direction, taking in energy (in the form of high-energy electrons) and building up larger molecules from smaller ones. This is called the reductive mode of the cycle. If an organism has access to high-energy electrons like those produced by geochemical processes, in fact, it can thrive with the cycle exclusively in the reductive mode, having no use for the oxidative mode at all. One way to think about the two modes of the cycle is this: In the oxidative mode, the input is an organic molecule, and the output is chemical energy, carbon dioxide and water. In the reductive mode, the input is chemical energy, carbon dioxide and water, and the output is a more complex molecule.

This must have been the way the cycle operated on the early Earth, because molecular oxygen was not available primordially to support the oxidative mode, and because we see it operating this way today in some anaerobic organisms that seem to have preserved this aspect of the biochemistry of their ancestors. In the reductive mode, the cycle provides a way for high-energy electrons to flow down the chemical hill. It is similar to the acetate reaction shown earlier, which is thermodynamically feasible but very slow, but with the addition of a network of small molecules—the reductive citric acid cycle—



acting to mediate and speed up the reaction. On biochemical and thermodynamic grounds, then, the reductive citric acid cycle (or some simpler precursor) would be a good candidate for the threshold of early life—the point where the pond of high-potential water is breached and the downhill pathway is etched out. The slow uncatalyzed conversion of carbon dioxide and hydrogen into acetate and water, shown earlier, occurs efficiently as the energy and reactants enter a primordial network of reactions like the modern-day reductive citric acid cycle.

In the metabolic maps of all modern organisms, the small molecules and reactions of the citric acid cycle are the starting point of every biosynthetic pathway—all roads lead from the citric acid cycle. However, in some organisms the reactions do not form a closed—cyclic—reaction sequence. For that reason, even among researchers convinced that these reactions are vestiges of the first metabolism, debate remains over whether the very first metabolic foothold was a cycle. However, because only cycles can act as self-amplifying channels, and because in organisms not running the closed cycle, sophisticated compensating adaptations are required, we consider a primordial reductive citric acid cycle the most likely route from geochemistry to life—the rivulet that formed at the top of the energy hill, through which the pond of energy began its thermodynamic escape. We then ask how, from this simple beginning, could the complexity we see in the modern cell arise. The first thing to notice is that, taken by itself, the cycle

captures only part of the energy in the carbon dioxide and hydrogen that constitute its input. In transforming the carbon dioxide to acetate, for example, the cycle harvests only about a third of the energy available in the electrons. Even in the deep core of metabolism, however, we do not see the cycle in isolation. Its lowest-energy molecule, acetate, is the starting point for other pathways that make the essential oils used in cell membranes, harvesting another third of the electron energy. Further reactions, such as those that generate methane, can capture the remaining available energy, though methane is a gas and therefore a waste product, unlike the earlier molecules in the pathway, which are constituents of biomass.

Selection Begins

We note that there is a fundamental difference between the way chemical reaction systems could have operated before the appearance of the first self-replicating molecules and the way they operate now that self-replicating systems have developed. In the beginning, the only potential source of order would have been networks of chemical reactions operating according to the laws of chemistry and physics. After molecules appeared that could replicate more or less independently, such as RNA, however, evolution could have proceeded according to the rules of natural selection, with the success of subsequent generations dependent on adaptive properties. Exactly when and how this transition occurred remains an open question debated by researchers, but the fact that it did occur is plain. Another way of saying this is that before the appearance of the first self-replicating molecules or assemblages of molecules (and, again, we have to emphasize that these may or may not have been inside cells), what mattered was the persistence of the chemical network; after such a system appeared, natural selection took on its more familiar form of selection among rival reproducing “individuals.”

Once natural selection began, systems with slightly different chemistry would appear on the scene through random accident. For example, acetate can be used in two ways to make oily molecules, and the major domains of life divide, in part, according to which class they make and how they use them. Methane production purely for energetic purposes may have been primordial, or it may have been coupled to metabo-

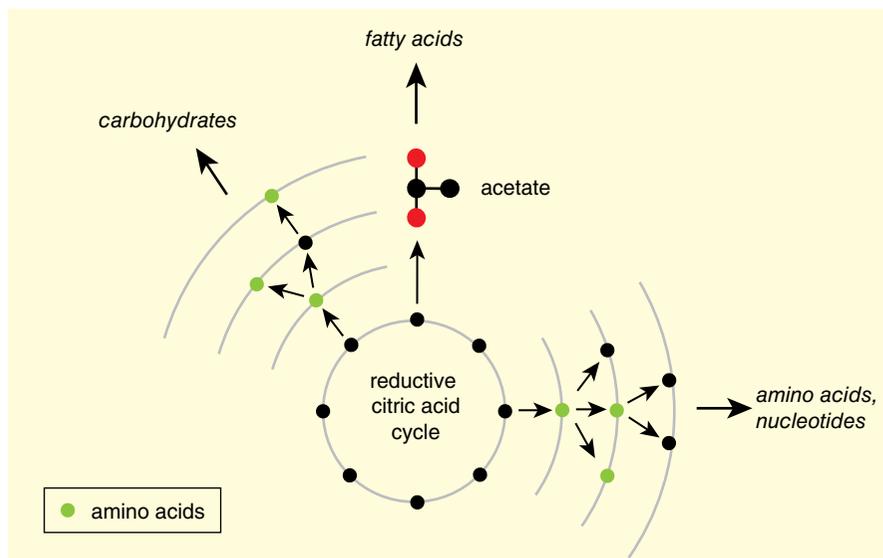


Figure 4. The famous metabolic map on page 209 was recently reworked to feature metabolism of chemoautotrophs—organisms that derive energy from inorganic sources and synthesize all necessary compounds from CO_2 . Arraying the reactions in concentric rings reveals a short route to essential biological products.

lism in a later, more complicated age (another topic of serious debate among researchers investigating the deepest branches of the tree of life).

The important pattern to appreciate is that the primordial cycle provides the stability and starting materials that make an age of selection possible. We think it was at the transition to this stage that geochemistry began to take on the features of replication and selection

recognized by Darwin as distinctive of life. After such an age has begun, it can maintain the complexity and diversity needed to explore for refinements—in efficiency, in adaptation to the geological environment or in specialized division of labor within communal systems. The same pattern repeated itself when the environment was changed by the accumulation of a destructive toxin—oxygen—that was produced by pri-

mordial organisms as a waste product. As they adapted, organisms did not abandon the reductive citric acid cycle, which we believe was the unique foundation for biosynthesis. Instead they acquired the ability to run the cycle in reverse, extracting energy from the breakdown of molecules similar to those the cycle formerly produced.

The role of the citric acid cycle as a foundation for complexity applies not only to subsequent adaptation by organisms under selection; it can be seen even within the chemical structure of the metabolic core itself. A particularly powerful way to make this point is to rework the schematic chart of current metabolism first developed by Nicholson. The original Nicholson chart was developed to elaborate human metabolism and was gradually expanded to incorporate the complex webs of chemistry on which humans depend. Recently, one of us (Morowitz) and Vijay Srinivasan used evidence from microbiology to distill the Nicholson chart, with its complex modules and domains of metabolism, down to a minimal common core, the necessary and sufficient network of reactions to make a living system. Within this core chart, which will be published soon, we arrayed pathways as layers built around citric acid cycle precursors. A fragment of that detailed chart is shown in Figure 4. The innermost layer consists of molecules that can be built from cycle intermediates with one chemical reaction, the next layer consists of those that can be built with two reactions, and so on. (Once you get past the first few layers, the counting becomes ambiguous, as the reactions often involve molecules that were themselves the products of layers farther in).

From this layered structure we believe we can see the chemical cascade that comprised the earliest steps in the evolution of life.

The primordial core chart is simpler than the elaborate chart made by combining organisms today, but it is not much simpler biosynthetically. It contains the major modules for sugars, oils and amino and nucleic acids, and we have proposed that it was—at least in broad outline—the agency of chemical selection in an era that preceded natural selection on distinguishable organisms.

If this notion turns out to be true, it will have important implications for a deep philosophical question: whether we should understand the history of life in terms of the working out of pre-

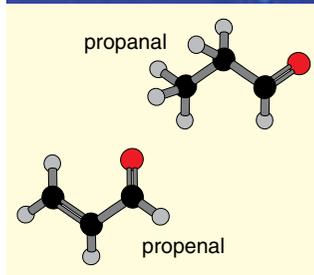
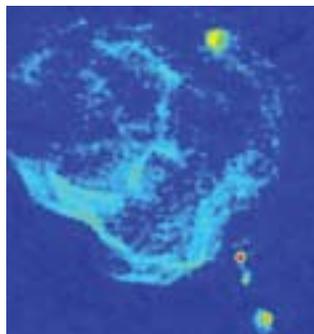


Figure 5. PRIMOS—the Prebiotic Interstellar Molecule Survey—has focused on Sagittarius B2(N), a cloud near the center of our galaxy (left, radio telescope image top right). The radio footprints of many organic molecules have been detected there (bottom right). Where aqueous carbon chemistry occurs, is metabolism far behind? (Image (left) courtesy of NASA and The Hubble Heritage Team (STScI/AURA); Image (top right) courtesy of Gaume, R., et al. 1995. *Astrophysics Journal* 449:663, reprinted by permission of the American Astronomical Society.)

dictable physical principles or of the agency of chance. We are, in fact, arguing that life will appear on any planet that reproduces the environmental and geological conditions that appeared on the early Earth, and that it will appear in order to solve precisely the sort of "stranded electron" problem discussed above. The currently popular view that complex life was something of a frozen accident was set forth in Jacques Monod's classic book *Chance and Necessity* (1970). We, of course, are arguing the opposite, if only for a significant part of basic chemical architecture. (It is important to appreciate that Monod studied regulatory systems, and in the domain of his expertise, we recognize the importance of accident, though we believe he advocated it too broadly.) It has not escaped our notice that the mechanism we are postulating immediately suggests that life is widespread in the universe, and can be expected to develop on any planet whose chemistry resembles that of the early Earth.

The view of life originating as a network of simple chemical reactions will require a lot of testing before it is adopted by the scientific community. We identify two areas where research is being pursued: the development of the theory of nonequilibrium statistical mechanics and the experimental pursuit of those first nonenzymatic chemical reactions that led to modern life.

On the theoretical side, we have to start with the realization that if we apply standard equilibrium thermodynamics to living systems, we arrive at something of a paradox. Living systems possess low entropy, which makes them very improbable from the equilibrium thermodynamic viewpoint. From the point of view of theoretical physics, the basic problem is that classical thermodynamics has only been well developed for systems in equilibrium—systems that do not change over time—or that change only by moving through successive, infinitesimally different equilibrium states. What is needed, therefore, is an extension of ordinary thermodynamics so that it can apply to systems maintained far from equilibrium by the flow of energy.

One promising approach was first suggested by E. T. Jaynes in the mid-20th century. He recognized that information (and hence entropy) is associated not just with states but with whole histories of change, which can include channel flows of the sort we have been discussing. Technically, one cannot talk about the entropy

"of a state" if the state depends for its context on a process of change; only the entropy of the whole process is expected to be maximized. To return to our pond on the hill, there is not a separate entropy of the pond, except as an approximation. Rather, there is an entropy of paths of change that include pond, channel, construction and relaxation. When such a formulation is analyzed for a simple system, the establishment of a channel can be seen as a phase transition, similar to the freezing of an ice cube or, to use a more precise mathematical analogy, the formation of a magnet from molten iron. (In the latter case, the phase transition occurs as the metal cools when the atomic dipole magnets line up in the same direction—paradoxically, a more ordered state). The full entropy of the process will be maximized in the system, even though the approximate entropy associated with the "state" of the channel may not be, thereby eliminating the paradox.

Current research into this foundational question now centers on the fact that the chemical substrate of living systems is much more complex than that of simple physical systems that have been examined so far. One important new direction of research involves the development of small-molecule catalysts in increasingly complex cooperative networks. The hope is that when a full theory is available, we will see the formation of life as an inevitable outcome of basic thermodynamics, like the freezing of ice cubes or the formation of magnets.

On the experimental side, some researchers, such as George Cody at the Carnegie Institution of Washington, D.C., are trying to work out the basic rules of organic chemistry for exotic environments that might have been relevant to the origin of life. Cody, for example, has worked on unraveling organic interactions at the kinds of temperatures and pressures that obtain at deep ocean vents. Mike Russell at the Jet Propulsion Laboratory in Pasadena, California, (author of "First Life," January–February 2006) is building a large chamber to model the geochemistry of those environments. Shelley Copley at the University of Colorado at Boulder has been sorting out the intermediate chemistry leading to the current nucleic acid–protein system of genetic coding, with an eye toward resolving the chicken-and-egg problem. These experiments represent a major paradigm shift from the top-down control envisioned in RNA World scenarios. Rather than

supposing that a few large RNA molecules control the adaptation of a passive small-molecule reaction network, Copley supposes that whole networks of intermediate molecules support each other on the path toward complexity. In this experimental setting, networks of small and randomly synthesized amino acids and single RNA units aid each others' formation, assembly into strings and evolution of catalytic capacity. Both types of molecules grow long together. Complexity, adaptation and control are distributed in such networks, rather than concentrated in one molecular species or reaction type. Distributed control is likely to be a central paradigm in the development of Metabolism First as a viable theory. We eagerly anticipate more experimental efforts like these to explore the many facets of small-molecule system organization.

In a larger sense, however, the future of the experimental program associated with the Metabolism First philosophy is tied to the development of the appropriate theory, guided by experimental results. The hope is that the interplay of theory and experiment, so familiar to historians of science, will produce a theory that illuminates the physical principles that led to the development of life and, hence, give us the ability to re-create life in our laboratories.

Assuming the experimental and theoretical programs outlined above work out well, our picture of life as a robust, inevitable outcome of certain geochemical processes will be on firm footing. Who knows? Maybe then someone will write a book titled *Necessity, Not Chance*.

Bibliography

- Morowitz, H. J. 1999. A theory of biochemical organization, metabolic pathways, and evolution. *Complexity* 4:39–53
- Smith, E., and H. J. Morowitz. 2004. Universality in intermediary metabolism. *Proceedings of the National Academy of Sciences of the U.S.A.* 101:13168–13173.
- Morowitz, H. J., and E. Smith. 2007. Energy flow and the organization of life. *Complexity* 13:51–59.
- Srinivasan, V., and H. J. Morowitz. 2009. The canonical network of autotrophic intermediary metabolism. *Biological Bulletin*. In Press.

For relevant Web links, consult this issue of *American Scientist Online*:

<http://www.americanscientist.org/issues/id.78/past.aspx>