

AMINO — The Immortal Thread

a journey through deep time, told by the molecules that remember

AICE Library · amino-the-complete-book

14 chapters · single-language edition

Every number, name and year in this volume is traceable to a cited source.

01. The First Choice

Four billion, two hundred million years ago, give or take — 4.2 Gya, in the shorthand this book will keep using, where Gya means giga-annum, a billion years before now — something on Earth is already alive. The team that reconstructed its chemistry favors a hydrothermal vent or a hot spring as its home: cracked rock, moving water, heat. This population already breathes by fixing carbon dioxide the way an anaerobic acetogen does, an organism that lives by turning carbon dioxide into acetate, through a chemical assembly line called the Wood-Ljungdahl pathway. Every protein its cells make is already built from one hand of amino acid, and only one.

That one-sidedness has a name: chirality, borrowed from the Greek word for hand. An object is chiral when it cannot be laid on top of its own mirror image and matched up, however it is turned. A left hand and a right hand share the same five fingers in the same order, and no rotation ever seats one inside the other's glove. Molecules can be chiral the same way: identical atoms, identical bonds, arranged as two mirror-image versions that no amount of twisting in space will superimpose. Chemists label the two versions L and D. Until 1848, nobody had proven that such mirror-twin molecules were real, physically separable things, rather than a mathematician's trick.

The proof came from a wine cask. Sodium ammonium paratartrate is a crystalline byproduct of winemaking, and by the 1840s chemists already knew it did something strange: dissolved in water, some batches twisted a beam of polarized light as it passed through, and other batches, chemically identical by every test then available, did not. In Paris in 1848, a young chemist named Louis Pasteur sat at a microscope with a pair of tweezers and sorted a heap of the inactive form, one tiny crystal at a time, by the tilt of facets too small to see unaided — some tilted one way, hemihedral to the right, others hemihedral to the left, the mirror opposite.

Then he dissolved the two piles separately and passed each solution through Jean-Baptiste Biot's polarimeter, an instrument built to measure exactly how far polarized light twists on its way through a liquid. In Pasteur's own words, preserved in a 2021 retrospective on the discovery: "I carefully separated the crystal hemihedral to the right and the crystals hemihedral to the left; I observed separately their solutions in Biot's polarizing apparatus, and I saw... the crystals hemihedral to the right deviated the plane of polarization to the right, those hemihedral to the left deviated it to the left." The two piles turned out to be one substance after all — a fifty-fifty mixture of two true mirror-image forms of tartaric acid, sorted apart by hand for the first time in history. Handedness was real. It could be separated, dissolved, and measured.

Life had already made this choice, long before Pasteur found a way to see it. Nearly every protein in nearly every organism alive today is built from L-amino acids, the left-handed form, strung together by the cell's own machinery. Nearly every strand of DNA and RNA rests on a backbone of D-sugars — D-deoxyribose in DNA, D-ribose in RNA — the right-handed form. Nothing in the underlying chemistry obviously demands this split; both hands are equally stable, equally easy to build in a laboratory flask. Life picked one hand for its proteins and the opposite hand for its genetic backbone, and at the level of its core machinery, it has never gone back.

That detail — opposite hands for the two great families of biological molecule — is itself an open question, not a solved one. Any single outside force strong enough to tip chemistry toward one handedness — circularly polarized starlight, a chiral mineral surface acting as a sorting table, even the faint leftward lean built into the weak nuclear force — should be expected to tip everything the same way, not opposite ways for proteins and for sugars. Origin-of-life researchers, writing as recently as 2025 and 2026, still say plainly that there is no widely accepted explanation for why the two major biopolymers ended up on opposite sides of the mirror. It is one of the honest gaps sitting at the center of this whole story.

The rule is strict enough that its exceptions are worth naming one at a time, starting with bacteria. More than ten distinct D-amino acids, chiefly D-alanine and D-glutamate, get built directly into the tough mesh of the bacterial cell wall — peptidoglycan — usually in the sequence L-alanine, D-glutamine, L-lysine (or a related molecule called meso-diaminopimelate), D-alanine, D-alanine. *Escherichia coli* and *Staphylococcus aureus* do this today, along with essentially every walled bacterium alive, while archaea and the plant and fungal cells that build walls of their own do not. This reads as a bacterial invention, layered on top of an already-settled handedness, rather than a surviving fragment from before handedness was decided.

A stranger exception belongs to a mammal. Platypus venom carries a dedicated enzyme, roughly 50 to 60 kilodaltons in size, whose entire job is to seize specific amino acids inside the venom's defensin-like peptides and natriuretic peptides — proteins first assembled the ordinary, left-handed way — and flip them to the right-handed form, recognizing a short signature sequence at one end of the chain. Reported in 2006, it was the first such isomerase, an enzyme that converts a molecule into its own mirror-image twin, ever documented in any mammal. The identical trick — build it left-handed, then flip it — evolved again independently in cone-snail venom and in frog skin secretions: three separate lineages, mollusk, amphibian, and monotreme, arriving at the same late workaround long after the original handedness was already fixed.

Exceptions this rare, this late, and this visibly bolted onto standard chemistry make the standard chemistry more striking, not less. Nothing in the physics of an amino acid or a sugar requires the left hand over the right; a coin has no reason to prefer heads. And yet nearly all of life, across billions of years and every branch it has ever grown, kept landing on the same side. A shared, arbitrary choice, preserved this faithfully through so much time and so much divergence, is exactly the kind of signature a family resemblance leaves behind: a clue that one ancestral population handed the choice down, rather than every lineage separately tossing the same coin and happening, again and again, to match.

One Ancestor, or Many Pushed the Same Way

A shared choice is not, on its own, proof of a shared origin. A 2025 review in PNAS lays out the honest counter-argument: an external chiral bias — the same circularly polarized starlight, the same parity-violating weak force, the same kind of chiral mineral surface named above — could in principle have leaned on more than one independent origin-of-life event and tipped every one of them toward the same hand, the way a strong wind can land many separately thrown coins the same way without the throws having anything to do with each other. Shared handedness is consistent with one first population. It is equally consistent with several, pushed the same direction by the same outside hand. Chirality alone does not settle which.

What can be dated, with real error bars attached, is the population that already had this chemistry locked in. A 2024 study in *Nature Ecology & Evolution*, led by Moody, Álvarez-Carretero, Mahendrarajah, and Donoghue, puts the last universal common ancestor of all life alive today — LUCA — at approximately 4.2 billion years old: a 95 percent credible interval of 4.09 to 4.33 Gya under the team's main relaxed molecular-clock model, and a full span of 3.94 to 4.52 Gya once every model they compared is counted. The method rests on five pairs of paralogous genes — genes born from a duplication inside one ancient genome, before LUCA existed, then inherited down two separate lines ever after — among them the two subunits of ATP synthase and the two elongation factors, Tu and G, calibrated against thirteen fossil and geochemical tie-points stretching from the Moon-forming impact, 4,510 million years ago, to a 2,954-million-year-old oxygen signature preserved in South Africa's Mozaan Group.

That figure revises an earlier reconstruction: Weiss and colleagues, writing in *Nature Microbiology* in 2016, described a thermophilic, nitrogen-fixing LUCA built from somewhat different evidence. The newer paper also argues that a ceiling some older studies leaned on — a supposed heavy bombardment of the early Earth by asteroids and comets, around 3.7 to 3.9 Gya — is not a safe upper bound, since the

timing, intensity, and even the reality of that bombardment episode is itself disputed in the lunar-impact record. LUCA's date is not an inscription. It is the current best answer, already revised once, open to being revised again.

By the time that population existed, whatever it looked like and wherever it lived, its cells were already reading something: a set of instructions, written in a molecular alphabet, translated then the same way its every descendant translates it now. The choice of hand came first. The choice of letters came next.

Four billion years, one alphabet

A timeline of the molecules that remember. Compressed toward the top; not to linear scale.



Overtuned if: a lineage of life — ancient, modern, or reconstructed in a laboratory — is shown to run on D-amino acids and L-sugars as its heritable core, rather than as a late, isolated exception, which would mean Earth's uniform handedness reflects one external push acting on separate origins rather than one shared ancestor.

Sources

- **2021** · Vantomme, "Pasteur and chirality: A story of how serendipity favors the prepared minds" — [Chirality \(Wiley\)](#)
- **2016** · "Why are all proteins 'left-handed'? New theory could solve origin of life mystery" — [Science/AAAS](#)
- **2011** · "Emerging knowledge of regulatory roles of d-amino acids in bacteria" — [Cellular and Molecular Life Sciences \(Springer\)](#)
- **2006** · Torres et al., "Mammalian L-to-D-amino-acid-residue isomerase from platypus venom" — [FEBS Letters](#)
- **2025** · "Life's homochirality: Across a prebiotic network" — [PNAS 122](#)
- **2024** · Moody et al., "The nature of the last universal common ancestor and its impact on the early Earth system" — [Nature Ecology & Evolution](#)

02. Four Letters

Something alive, still close to 4.2 billion years old, is already reading. Its alphabet has exactly four letters: adenine and thymine, guanine and cytosine, in the double-stranded archive called DNA, or adenine and uracil standing in for thymine, in the single-stranded working copy called RNA. The letters do not pair at random. Adenine always finds thymine, or uracil, and guanine always finds cytosine, the two strands running alongside each other like a zipper whose teeth fit only one specific partner apiece. That fixed pairing is the whole reason a cell can copy its archive at all: split the zipper down the middle, and each half already spells out exactly what its missing partner must be.

Four letters would be a thin alphabet for spelling out a body, except that the cell never reads one letter at a time. It reads in words of exactly three letters, called codons, run end to end with no punctuation and no gaps between them — an arrangement the Nobel committee's own 1989 press release described in almost exactly these terms: "only four such letters... arranged as a long sentence of three-letter words." Three letters, drawn from an alphabet of four, is not an arbitrary word length. It is the shortest word length able to give the cell enough distinct spellings to name every amino acid it needs, with room to spare.

The arithmetic is worth doing in full, because it turns out to run the whole living world. Four possible letters at the first position, four at the second, four at the third: four times four times four, sixty-four possible three-letter codons, no more and no fewer. Of those sixty-four, sixty-one specify one of the twenty amino acids that proteins are built from. The remaining three — UAA, UAG, and UGA — carry no amino acid at all; they are stop signs, telling the cell's protein-building machinery exactly where a protein ends. One further codon, AUG, pulls double duty: it is both the standard start signal and the codon for the amino acid methionine, a capital letter that also happens to spell a word on its own. This entire structure was cracked experimentally, codon by codon, in the early-to-mid 1960s, chiefly by Marshall Nirenberg, Heinrich Matthaei, and Har

Gobind Khorana, and it is taught identically, without local variation, in essentially every genetics classroom since.

Sixty-four codons for twenty amino acids and one stop signal means most amino acids answer to more than one spelling, a property biologists call degeneracy, or redundancy. The pattern has its own precise arithmetic: five amino acids — proline, alanine, threonine, valine, and glycine — each get four different codons, a set biologists call a quartet. Nine amino acids get exactly two codons each, a doublet. Three amino acids — leucine, arginine, and serine — get six codons each, a sextet, the richest spelling of all. One amino acid, isoleucine, gets three codons, a triplet. Two amino acids, methionine and tryptophan, get exactly one codon each, a singlet. Multiply it out and it closes perfectly: five quartets plus nine doublets plus three sextets plus one triplet plus two singlets comes to twenty amino acids, spelled across sixty-one codons, with nothing left over and nothing missing.

That redundancy is not sloppiness. It is insurance, built directly into the alphabet. Because so many amino acids have backup spellings, a large fraction of the copying errors that strike a gene's third codon letter change nothing at all about the resulting protein — insurance the cell gets for free, simply because of how the sixty-four spellings were handed out. The insurance runs deeper still: Francis Crick proposed in 1966 that the pairing between a codon and its matching anticodon, the three-letter tag carried by the transfer RNA molecule that ferries each amino acid into place, is strict at the first two letters but tolerant, or "wobbly," at the third. Because of that wobble, a single transfer RNA can often read more than one synonymous codon, which is why a cell gets by with roughly forty distinct transfer RNA types instead of needing a separate one for every one of the sixty-one sense codons.

The redundancy pattern turns out to do something more than buffer typos: it buffers meaning. In 1998, Stephen Freeland and Laurence Hurst measured the standard code against roughly one million computer-generated alternative codes, all built under the same basic rules, and scored each one for "mutational cost" — how much damage a single-

letter copying error tends to do to the resulting protein, given that codons differing by only one letter tend to specify chemically similar amino acids. The real, standard code beat more than 99.99 percent of the alternatives. Freeland extended the finding in 2000 with Knight, Landweber, and Hurst, arguing that this near-optimal arrangement was essentially locked in before LUCA existed. Not every researcher accepts that this reflects deliberate selection for error resistance: a 2010 paper in PNAS, "A Neutral Origin for Error Minimization in the Genetic Code," argues the same low-damage pattern could have emerged as a side effect of how codon assignments accumulated over time, through a process of reducing ambiguity, rather than through direct selection for minimizing errors. The debate remains open in the primary literature.

The Code's Oldest Argument

A second, older argument concerns not how well the code resists damage but why these particular letters ended up paired with these particular amino acids at all. Francis Crick proposed one answer in 1968: the "frozen accident." Once a translation system and a great many proteins already depended on a given set of codon meanings, reassigning even one codon would have been catastrophic to everything built on top of it, so the assignments locked in place by historical accident and simply never moved again. Carl Woese proposed a rival answer starting in 1965: the stereochemical theory, which holds that codon or anticodon sequences carry direct, if weak, chemical affinities for their own amino acids, so the code's shape reflects real molecular binding preferences rather than pure historical lock-in. Laboratory binding assays that Woese ran in 1966 and 1967, and again with Saxinger and Ponnamperna in 1971, found only what the record describes as "weak and relatively non-specific interactions" — evidence for the idea, but not strong evidence. Most current reviews, including a 2017 survey by Eugene Koonin, treat the honest answer as a blend of both: some early stereochemical shaping while the code was still forming, followed by frozen-accident lock-in once code and genome had grown too tightly interdependent to safely change.

Interdependent, though, has never meant unchangeable everywhere. A small number of living lineages run non-standard versions of the code today. Many animal mitochondrial genomes reassign some of their stop or sense codons; certain ciliates reassign UAA and UAG from stop signals to the amino acid glutamine; *Mycoplasma* bacteria reassign UGA from a stop signal to the amino acid tryptophan. Human mitochondria are one of the living examples, reassigning codons including AGA and AGG. Far from undermining the frozen-accident idea, these exceptions support it: they show up specifically in small, gene-poor genomes, like mitochondria, where few enough proteins depend on the affected codon that reassigning it is survivable rather than catastrophic — exactly the condition frozen-accident logic predicts should make change possible.

Step back from the mechanism entirely and look at the code as a channel carrying information, and a different kind of number appears. Each of the four bases, treated as one symbol in a four-symbol alphabet, carries a theoretical maximum of two bits of information — log base two of four, a mathematical identity, not a measurement. Real genomic DNA, measured empirically, carries somewhat less: approximately 1.9 bits per base, because real sequences are not perfectly random. Biased base composition and codon-usage bias both pull the realized information down below the theoretical ceiling, and non-coding stretches of the genome measure lower still than coding stretches do. Two bits, in other words, is the speed limit; 1.9 is closer to what life actually drives.

Set beside any human information system, the comparison is almost unfair. The standard code, cracked in the 1960s and barely modified since except in a handful of small, late, low-cost corners, is still running, essentially unchanged, in a bacterium and in every cell of the organism reading about it now. It is very likely the oldest continuously operating data format on the planet, never rewritten from scratch, only patched at the margins, carrying the same twenty-amino-acid, sixty-four-codon logic across more than four billion years of unbroken use.

An alphabet, on its own, is inert. It needs something to read it, and something to act on what it reads — and by the time LUCA's genome held

this code, both of those already existed, doing a job neither DNA nor protein can do alone.

The oldest data format on Earth

Four letters, three-letter words, twenty meanings — the same code in a bacterium and in a reader.

1 · Four letters, paired



2 bonds



3 bonds

A always pairs with T
(U in RNA); G always with C.
Each strand is the recipe
for rebuilding the other.

2 · Read three at a time — a codon



one codon = one 3-letter word

$$4 \times 4 \times 4 = 64$$

64 possible words from just four letters.

Two letters (16) would be too few for 20 amino acids;
three is the smallest word that works.

3 · 64 words → 20 amino acids + STOP

64 codons

every 3-letter word



20 amino acids

+ a STOP signal

More words than meanings,
so meanings repeat.

Leucine is spelled six ways:

U U A

U U G

C U U

C U C

C U A

C U G

The spare spellings are not waste: a mis-copied third letter often still means the same amino acid,
so the redundancy quietly absorbs copying errors. The code protects its own meaning.

4 · Counting it in bits

One letter

2 bits

$\log_2 4$ — a choice of one out of four

One codon

6 bits

three letters, $2 + 2 + 2 = 6$ states

To name 1 of 20

~4.32 bits

$\log_2 20$ — what the meaning actually needs

The difference

~1.7 bits

spent on redundancy, banked as error protection

The genetic code, counted out.

Overtaken if: a large, ancient, gene-rich genome — not a small, reduced one like a mitochondrion — is found to reassign a sense

codon's meaning at negligible fitness cost, showing the code's near-universality reflects easy interchangeability rather than a catastrophic, frozen lock-in.

Sources

- **2020s (open textbook, accessed 2026)** · "An introduction to the genetic code," Chromosomes, Genes, and Traits: An Introduction to Genetics — [Rotel/Pressbooks](#)
- **1989** · Nobel Prize in Chemistry Press Release, "Ribonucleic acid (RNA) – a biomolecule of many functions" — [NobelPrize.org](#)
- **2011** · "The multiplet structure of the genetic code, in one and small number" — [arXiv:1101.2983](#)
- **1998** · Freeland & Hurst, "The Case for an Error Minimizing Standard Genetic Code" — [Journal of Molecular Evolution \(Springer\)](#); extended by **2000** · Freeland, Knight, Landweber & Hurst, "Early Fixation of an Optimal Genetic Code" — [Molecular Biology and Evolution \(Oxford Academic\)](#)
- **2017** · Koonin, "Frozen Accident Pushing 50: Stereochemistry, Expansion, and Chance in the Evolution of the Genetic Code" — [Life \(MDPI\)](#) 7(2):22
- **2012** · "On the Entropy of DNA: Algorithms and Measurements Based on Memory" — [Wharton School faculty paper](#)

03. The Molecule That Copied Itself

Somewhere in the interval between roughly 4.5 billion years ago, when the Earth-Moon system had only just finished forming, and roughly 4.2 billion years ago, when LUCA's population already carried a working genetic code, something existed that had not yet split its two jobs apart. The 2024 study that dates LUCA calls this interval "surprisingly short." Inside it, one kind of molecule appears to have served as both the message and the machine reading it — a single substance holding heredity and chemistry at once, before life divided that work between DNA, the archive, and protein, the machinery.

That molecule is RNA, and the idea that an entire era of early life ran on RNA alone has a name and a birth date: the "RNA world," coined by Walter Gilbert of Harvard University's Biological Laboratories in a one-paragraph commentary in *Nature*, volume 319, page 618, in the issue dated 20 February 1986. Gilbert was the first to propose explicitly that the earliest living systems were built entirely from RNA, carrying out both information storage and catalysis — the triggering and speeding-up of chemical reactions — in one molecule, roles later split between DNA and protein. The commentary is short, but by one count from *Nature*'s own metrics, it has since been cited 2,623 times.

Gilbert had a year-old discovery to build his idea on. Two researchers, working independently, had just shown that RNA could act as a genuine biological catalyst — a ribozyme — rather than only a carrier of genetic information. Thomas Cech, at the University of Colorado, Boulder, showed in 1982 that an unprocessed RNA molecule from the ciliate *Tetrahymena thermophila* could splice itself, cutting and rejoining its own sequence with no protein present anywhere in the reaction. Sidney Altman, at Yale University, built his case in stages: from a 1978 finding that the bacterial enzyme RNase P lost its function when its RNA component was removed and regained it when the RNA was put back, through to a fuller 1983 demonstration that the RNA component alone, unassisted, catalyzes the cutting reaction RNase P is known for. Cech and Altman

shared the 1989 Nobel Prize in Chemistry "for their discovery of catalytic properties of RNA." By the time the prize was announced, "close to a hundred" such ribozymes were already known.

Neither discovery is a museum piece. RNase P, which processes transfer RNA, runs today in essentially every bacterium, archaeon, and eukaryote alive. Cech's self-splicing introns are still found in *Tetrahymena* and in the genes of many other organisms besides.

The single strongest piece of evidence for an RNA-run past, though, sits at the exact center of the machine every living cell still depends on to build proteins at all: the ribosome. In 2000, Nissen, Hansen, Ban, Moore, and Steitz published the atomic-resolution structure of the large ribosomal subunit from the archaeon *Haloarcula marismortui* in *Science*, volume 289, pages 920 to 930, and found that the peptidyl transferase center — the exact site where a growing protein chain receives its next amino acid — is built entirely from ribosomal RNA. No protein side-chain atom sits closer than about 18 angstroms to the reaction itself, a distance far too great to take part; both substrate molecules touch only conserved RNA belonging to a section of the ribosome called domain V. Every ribosome alive today, in every domain of life, still runs this same RNA-only reaction chamber. It may be the single most universal fossil the RNA world left behind, because it sits at the very center of biology's most essential machine rather than in some peripheral, vestigial corner of it.

RNA's capacity to copy itself, rather than merely catalyze a single reaction, was demonstrated directly in the laboratory in 2009. Tracey Lincoln and Gerald Joyce, at the Scripps Research Institute, engineered a pair of RNA enzymes that catalyze each other's synthesis from four short building-block molecules, so that each enzyme manufactures its own partner. The result, published in *Science*, volume 323, pages 1229 to 1232, was self-sustained, exponential amplification with a doubling time of about one hour, continuing indefinitely as long as fresh raw material kept arriving. When several variant versions of the replicators were set to compete for one shared pool of material, new recombinant replicators appeared during the run and came to dominate the population —

Darwinian evolution, acting on a purely chemical system, with no protein involved anywhere. The honest limit on the experiment is one its own authors would recognize: whether this particular, carefully engineered pair resembles how a natural, prebiotic replicator could first have arisen on its own, rather than requiring the kind of deliberate, iterative human design that built these two, is a fair question the result cannot answer by itself.

Where the Raw Material Came From

None of this chemistry works without amino acids and nucleotides already lying around to be strung together, and the question of where those came from has its own experimental history, starting in 1953. Stanley Miller, working with Harold Urey at the University of Chicago, sealed a mixture of methane, ammonia, hydrogen, and water vapor inside a flask, a guess at early Earth's atmosphere, and ran an electric spark through it continuously for about a week, standing in for lightning. Nineteen organic compounds turned up afterward, including the amino acids glycine, alpha-alanine, and beta-alanine, built from nothing but those four inorganic starting materials and a spark. The manuscript reached *Science* in February 1953.

Fifty-five years later, the same source material yielded more. In 2008, a team led by Jeffrey Bada rediscovered sealed sample vials from an experiment Miller himself had run in the 1950s but never published, using a modified "volcanic apparatus" built to simulate a water-vapor-rich volcanic eruption laced with lightning. Reanalyzed with instruments Miller never had access to — modern high-performance liquid chromatography and mass spectrometry — those decades-old vials gave up 22 different amino acids and 5 amines, several never before identified in any Miller-type experiment, a notably richer haul than the original published setup produced. Published in *Science* on 17 October 2008, the reanalysis argues that lightning inside volcanic eruption plumes, a genuine feature of the young, volcanically active Earth, may have been an especially productive corner of the planet for assembling the raw

material of life. The chemistry itself is not disputed. What is disputed is how closely Miller's assumed atmosphere, strongly reducing and rich in methane and ammonia, matched the real early atmosphere, which most current atmospheric science now judges was probably closer to a neutral mix rich in carbon dioxide and nitrogen instead.

A wholly different kind of evidence for the same chemistry arrives from outside the planet. The Murchison meteorite fell in Australia in 1969, and it carries a small but real excess of one mirror-image form over the other in the amino acid isovaline, a compound essentially absent from terrestrial biology, which is exactly why it can be trusted as an uncontaminated tracer of extraterrestrial chemistry. Measured excesses of the left-handed form range from none at all up to between 18.5 and 20.5 percent, depending on the specific stone and analytical method used — the largest such excess reported in any meteorite to date. Pizzarello's 2003 measurements, published in *Geochimica et Cosmochimica Acta*, found a range of 0 to 15.2 percent; a 2009 study in *PNAS* by Glavin and Dworkin measured 18.5 percent, plus or minus 2.6, specifically in Murchison, and up to 20.5 percent across a wider survey of carbonaceous chondrites, a class of stony meteorite rich in carbon compounds.

Murchison's parent body dates to the early Solar System, roughly 4.5 to 4.6 billion years ago, the standard formation age for meteorites of its class — meaning this chiral bias existed in Solar System rock before Earth's own homochirality was ever locked in. That timing is consistent with, though it does not prove, an extraterrestrial contribution to Earth's original tilt toward one hand. Glavin and Dworkin's own analysis adds an honest complication rather than resolving one: an excess as large as roughly 20.5 percent looks too big to explain by the leading proposed mechanism, irradiation by circularly polarized light, alone, and they propose aqueous alteration — chemical reworking by liquid water moving through the meteorite's parent body — as an additional or alternative process. The gap between mechanism and observation has not closed.

Between a spark that can write amino acids out of raw gas, a meteorite carrying a chiral bias older than Earth itself, and a laboratory

demonstration that RNA alone can copy itself, compete, and evolve, the outline of heredity's first moment comes into view: a molecule that copies itself, imperfectly, inside a pool already stocked with its own raw material. Imperfect copying, repeated, is the entire mechanical root of heredity — one generation resembling, but never perfectly repeating, the one before it.

Something that could copy itself, and occasionally copy itself wrong, already had everything it needed to start accumulating history rather than just chemistry. What it did not yet have was permanence: a lock strong enough to hold the arrangement it had found. That lock was still ahead.

Overtaken if: the chemistry preceding heredity is shown to have required peptides or a broader metabolic network working alongside RNA from the start, rather than a genuine stage in which RNA alone carried out both message and machine.

Sources

- **1986** · Gilbert, "Origin of life: The RNA world" — [Nature 319, 618](#)
- **2000** · Nissen, Hansen, Ban, Moore & Steitz, "The Structural Basis of Ribosome Activity in Peptide Bond Synthesis" — [Science 289\(5481\):920](#)
- **2009** · Lincoln & Joyce, "Self-Sustained Replication of an RNA Enzyme" — [Science 323\(5918\):1229-1232](#)
- **1953** · Miller-Urey spark-discharge experiment, NASA historical summary — [NASA NTRS technical report](#)
- **2008** · Johnson et al., "The Miller Volcanic Spark Discharge Experiment" — [Science 322\(5900\):404](#)
- **2009** · Glavin & Dworkin, "Enrichment of the amino acid L-isovaline by aqueous alteration on CI and CM meteorite parent bodies" — [PNAS 106\(14\)](#)

04. The Frozen Code

Four billion, two hundred million years ago, give or take a hundred million years either way, warm water seeps through a maze of mineral pores in the dark. No sunlight reaches this deep. This is the reconstructed home of the last common ancestor of everything alive today: not a fossil anyone has ever held, not a single cell anyone has ever seen, but a population inferred backward from the genes every living lineage still carries. Inside those pores, molecular machinery is already reading a dictionary of sixty-four three-letter words — codons, the basic units of the genetic code, each one a triplet of RNA or DNA letters that specifies one amino acid or a stop signal — and translating them into the proteins that build every one of its descendants. That dictionary has not changed since. It is still running, word for word, in every root, every reef, every hand that has ever turned a page.

How a dictionary this arbitrary gets adopted everywhere and then stops changing is itself a puzzle, and the man who gave the classic answer was Francis Crick, in 1968, in the *Journal of Molecular Biology*. Crick's idea has a plain mechanical logic: once enough proteins in a cell depend on one codon meaning one particular amino acid, changing that meaning would corrupt every protein already built on the old definition — like changing what a common word means after a million pages have already been printed using it the old way. Small tweaks that do not matter can still happen at the edges. Wholesale reassignment cannot, because it would break too much machinery at once to survive. Crick called this a frozen accident, and it remains the default explanation every rival theory has to argue against.

A frozen accident explains why the code stopped changing. It does not by itself explain why the code took the particular shape it froze into — why, for instance, similar codons tend to specify chemically similar amino acids rather than being scattered at random. Three later ideas compete to fill that gap, and none has driven the frozen-accident idea out of the

field; modern reviews tend to treat all four as layers rather than as rivals fighting for a single winner.

The first, proposed by J. Tze-Fei Wong in 1975 and reaffirmed in his own thirty-years-later review, argues that codon assignments track the biosynthetic pathways that build the amino acids themselves. Amino acids chemically derived from one another — aspartate turning into asparagine through a short chemical pathway, to take one example — tend to share related codons, as though the code expanded outward in step with the invention of each amino acid's manufacturing route, simple ones first, complex ones added later, the whole sequence finished before the last common ancestor existed. A later reanalysis in PNAS pushed back hard, arguing the pattern could be a side effect of amino acids simply resembling each other chemically rather than genuine evidence of biosynthetic history — a dispute filed under the blunt heading "fact or artifact." One later tally attempting to weigh the competing adaptive theories against each other put biosynthesis, error-minimization and chemical shape at a relative statistical influence of roughly forty million to four hundred to one — biosynthesis, on that reckoning, doing most of the explanatory work, though the comparison itself remains contested.

The second idea, from Stephen Freeland and Laurence Hurst in 1998, is easier to test and gave its own paper a memorable title: "The genetic code is one in a million." Freeland and Hurst built huge numbers of computer-generated alternative genetic codes and scored each one, including the real one, on how much damage a single misreading would do. Because codons for chemically similar amino acids sit near each other in the real table, a wrong letter tends to substitute a similar-shaped amino acid rather than a wildly different one, softening the damage of ordinary translation errors and mutations. The real code beat all but somewhere between one in ten thousand and one in a million randomly generated rivals, depending on exactly how errors were weighted. That is not the signature of an arbitrary table; it reads like a code pruned by natural selection before it ever froze.

The third idea reaches further back than biology itself. Michael Yarus, working with Jeff Widmann and Rob Knight, proposed that some codon assignments preserve a direct chemical handshake between an RNA triplet and the amino acid it encodes, no protein machinery required, just chemistry. Short RNA snippets fished out of test tubes purely for their ability to grip one particular amino acid, with no information about the genetic code built into the experiment, turned out to be enriched for that same amino acid's own codon or anticodon sequence far more often than chance would predict, in results published in 2009. The pattern held for seven of the twenty amino acids: arginine, tryptophan, histidine, tyrosine, isoleucine, leucine and phenylalanine. David Johnson and Lei Wang found an echo of the same chemical fingerprint in 2010, buried inside the structure of the modern ribosome itself. Seven out of twenty is a real signal, not the whole story, and coevolution theorists argue biosynthesis still explains more of the table overall.

The exceptions that prove the freeze

If the code had truly stayed fluid all the way to the present, variant versions should turn up scattered evenly across the tree of life. They do not. The NCBI's taxonomy resource, compiled by Andrzej Elzanowski and Jim Ostell and last updated in September 2024, catalogues thirty-three distinct named genetic-code tables: real, documented departures from the standard dictionary. Vertebrate mitochondria reassign the codons AGA and AGG from arginine to a stop signal, and AUA from isoleucine to methionine. *Mycoplasma* and *Spiroplasma* bacteria, along with some mold and protozoan mitochondria, read UGA, ordinarily a stop signal, as tryptophan instead. The fungal pathogens *Candida albicans* and *Candida auris* read CUG as serine rather than leucine. Ciliates such as *Oxytricha*, *Tetrahymena* and *Paramecium* read UAA and UAG as glutamine instead of stopping there at all.

Every one of these lives on a single narrow branch — one mitochondrial lineage, one group of ciliates, one or two fungi — never anywhere near the root of the tree. That pattern is the signature of late, local tinkering

on a table that had already frozen solid, not a surviving fragment of a code still finding its footing. Exactly how a lineage survives such a swap without poisoning every protein it owns is still debated, under names like codon capture and ambiguous intermediate, but that argument concerns the mechanics of a late edit. It does not reopen the question of whether the ancestral table itself was ever unfrozen.

That same universality is the strongest reason to think Bacteria, Archaea and Eukarya, the three domains into which all cellular life divides, share one common inheritance rather than three independent inventions that happened to converge. A code this arbitrary is not the kind of thing three separate lineages reinvent and land on identically by chance. The shared ancestor implied by that logic has a name: LUCA, the last universal common ancestor.

In 2016, Madeline Weiss, William Martin and colleagues at Heinrich Heine University Düsseldorf tried to reconstruct what LUCA actually was, not by guessing but by screening 6.1 million protein-coding genes drawn from sequenced microbial genomes, sorting them into 286,514 clusters of related genes, and tracing 355 of those families, barely one gene in a thousand, back to LUCA with statistical confidence. What falls out of that reconstruction is a physiology, not a portrait: an anaerobic organism living without oxygen, running on hydrogen gas, tolerant of high heat, fixing carbon dioxide into organic matter through a biochemical route called the Wood-Ljungdahl pathway, in a hydrothermal-vent-type habitat. Modern microbes called acetogens and methanogens still run close to that same core chemistry today, at vents, in swamps, in guts, wherever oxygen does not reach.

Weiss and Martin's paper reconstructs physiology, not a date. The date came in 2024, from Edmund Moody at the University of Bristol and colleagues, who built a molecular clock, a technique that reads the small accumulated differences between genes as a kind of stopwatch calibrated against independently known events, out of five separate sets of paralogous genes: pairs of genes within a single genome that arose when one ancestral gene duplicated and each copy then drifted on its own.

Cross-checked against the microbial fossil and isotope record, their estimate placed LUCA at roughly 4.2 billion years ago, with 95 percent confidence spanning 4.09 to 4.33 billion years, early enough to brush against the tail end of the Late Heavy Bombardment, the period when the inner solar system was still being pelted by leftover debris from its own formation. That date implies something uncomfortable: a genetically complex, ecosystem-embedded organism, not a bare minimal proto-cell, was already running almost as soon as the planet had cooled enough to host it. The method is not seriously disputed. The date it produces runs well earlier than the 3.5-to-3.8-billion-year window most older estimates favored, and that gap remains a live argument.

One correction matters more than any date, though: LUCA was very likely never a single organism, the way a family tree draws one dot at the top and everything descends from it. Nicolas Glansdorff, Ying Xu and Bernard Labedan argued this in 2008, comparing gene trees across sequenced genomes and finding that different gene families trace back to common-ancestor points that do not coincide at one shared moment. Moody's 2024 reconstruction leans the same way, describing LUCA as one member of a contemporaneous, interacting ecosystem rather than an isolated founder. The reason the ancestor points scatter is horizontal gene transfer: genes passed sideways, between unrelated cells living at the same time, rather than only straight down from parent to offspring. Picture LUCA not as a single primordial cell whose descendants later split into three domains, but as a loosely bounded, gene-swapping neighborhood of related cells, no single one of which was uniquely first. The evidence sits in plain view today: Bacteria and Archaea do not carry one identical inherited toolkit passed down whole. They carry different, only partly overlapping slices of it, the way cousins might each inherit a different subset of a shared grandparent's belongings rather than one matching set.

Whatever LUCA was, a population, a neighborhood, thousands of gene-trading cells sharing one already-frozen dictionary, it still kept that dictionary's records in the older, more fragile molecule. RNA had done the job since the RNA world, roughly four billion years earlier still. RNA is

not built to last, and the lineage's next move was to find something sturdier to hold the archive.

Overtured if: a genome basal enough to predate the split of Bacteria, Archaea and Eukarya were found running a fundamentally different, non-degenerate codon table — proof the ancestral code was still fluid at LUCA's own time rather than already frozen shut.

Sources

- **1968** · Crick, F.H., "The Origin of the Genetic Code" (reviewed in Koonin & Novozhilov, "Origin and Evolution of the Genetic Code") — [PMC5492144](#)
- **1975** · Wong, J.T., "A Co-Evolution Theory of the Genetic Code" — [PNAS 72\(5\):1909–1912](#)
- **2024** · Elzanowski, A. & Ostell, J. (comp.), "The Genetic Codes" — [NCBI Taxonomy resource](#)
- **2016** · Weiss, M.C. et al., "The physiology and habitat of the last universal common ancestor" — [Nature Microbiology 1:16116](#)
- **2024** · Moody, E.R.R. et al., "The nature of the last universal common ancestor and its impact on the early Earth system" — [Nature Ecology & Evolution](#)
- **2008** · Glansdorff, N., Xu, Y., Labedan, B., "The Last Universal Common Ancestor: emergence, constitution and genetic legacy of an elusive forerunner" — [Biology Direct 3:29, PMC2478661](#)

05. The Upgrade to DNA

A single strand of RNA sits in warm, faintly alkaline water, roughly four billion years ago, in the same window when the RNA world was still running and LUCA's population was assembling the frozen code from the last leg of this journey. At one point along its backbone, a small chemical group, a hydroxyl, one oxygen and one hydrogen sitting on the second carbon of the sugar ring, reaches around and attacks the very bond holding the backbone together. Nothing hostile does this. It is built into the molecule's own chemistry: give that hydroxyl group a little heat, a stray metal ion, or alkaline water, and it will cut its neighbor's throat. Every RNA molecule, then and now, carries this flaw as standard equipment.

DNA's fix is to remove the flaw at its source. DNA's sugar, deoxyribose, is missing exactly that one hydroxyl group, which is what the "deoxy" in its name means: literally, without that oxygen. No hydroxyl, no self-inflicted knife. Tomas Lindahl's classic 1993 review quantifies the payoff: DNA's backbone bonds are roughly two hundred times more stable than RNA's, measured at the neutral pH and physiological magnesium concentrations a living cell actually runs. That is not a trivial improvement. It is the difference between an archive that can hold its shape for a lifetime and one that cannot.

The logic that follows is simple and unforgiving. Any lineage able to move its master archive onto the sturdier molecule would suffer fewer spontaneous backbone breaks per generation, a standing chemical pressure favoring the switch the moment it became biochemically possible. DNA does not eliminate copying errors; it only slows the rate of spontaneous decay enough to buy time between them, time for repair systems to catch mistakes, time for a genome to grow larger without collapsing under its own damage. The gamble still shows up today in the handful of lineages that chose not to make the switch: influenza, coronaviruses and HIV all still carry their genomes as RNA, and all three

pay for it in mutation rates far higher than any cellular genome could survive carrying in its own chromosomes.

A second clue to the same transition hides in a single altered letter. DNA uses thymine, essentially uracil with a small methyl tag bolted on, where RNA uses plain uracil, and the reason reads like a forensic signature system. Cytosine, one of the four genetic letters, spontaneously degrades into uracil at a slow, steady background rate through a chemical process called deamination, in which the molecule loses a nitrogen-and-hydrogen group and turns into a different base entirely. If DNA used plain uracil the way RNA does, there would be no way to distinguish a correct original letter from that kind of damage. By making the real letter thymine instead, any uracil found sitting in a DNA strand is automatically recognizable as a mistake rather than a legitimate letter.

Cells exploit that signature with dedicated machinery. Four distinct human enzymes, UNG2, SMUG1, TDG and MBD4, each one a uracil-DNA glycosylase, a molecular proofreader whose only job is spotting uracil where thymine belongs, recognize and cut out that uracil, triggering a cut-patch-and-reseal cycle called base-excision repair. Hans Krokan and Magnar Bjørås's 2013 review calls this pathway evolutionarily highly conserved. No independent date is attached to that phrase, but conservation this deep, across such distantly related organisms, usually means old, plausibly tracing back to the transition itself.

A third clue sits in the machinery that builds DNA's raw materials, and it is the most telling of the three: there is no known biological route to make a DNA building block from scratch. The only known path runs through an enzyme called ribonucleotide reductase, which takes an RNA building block and strips off that same troublesome hydroxyl group, converting it directly into a DNA building block. Every dividing cell alive today still makes its DNA out of chemically modified RNA parts, every single time it divides. That is not a metaphor; it is current, ongoing biochemistry, and it is a fossil record hiding in plain sight. DNA synthesis is still mechanically parasitic on RNA chemistry, exactly as it would be if DNA arrived second.

The enzyme's own history reinforces the point. In 1997, Jordi Riera and colleagues cloned and sequenced this reductase from *Pyrococcus furiosus*, a heat-loving archaeon, and found it to be a class II, adenosylcobalamin-dependent enzyme carrying two intein insertions, a structurally distinct form that nonetheless shares strong sequence resemblance with the two other known classes of the same enzyme running in aerobic and strictly anaerobic organisms today. In 2015, Daniel Lundin and colleagues synthesized this into a single hypothetical ancestral enzyme, nicknamed the urRNR, that diversified into the three classes still dividing the tree of life: Class I in oxygen-breathing organisms, Class II in archaea and bacteria such as *P. furiosus*, Class III in strict anaerobes. Lundin's own team is candid that the true ancestral enzyme has never been identified, and calls it an open enigma whether an RNA-based enzyme, a ribozyme, doing a job usually left to proteins, could ever have managed this particular, radical-driven chemistry unassisted.

The dates bound each other. The RNA world is conventionally placed at roughly four billion years ago; Paul Higgs and Niles Lehman's 2015 review, treated in the field as the standard modern synthesis on the question, draws together comparative genomics, ribozyme-catalysis experiments and molecular-clock data to reach that figure, and Lehman's own companion editorial states it without hedging: the RNA world is "now some four billion years behind us." The switch to DNA-and-protein genomes is squeezed into the same narrow window, essentially complete by the time LUCA's population was already running its frozen code, roughly 4.2 billion years ago by the previous leg's dating. How long the RNA world itself lasted before that handover is still argued over, anywhere from perhaps a hundred million years to considerably longer.

Two molecules, two jobs

Once DNA existed as an option, life did not simply trade one molecule for the other. It split one job into two. DNA became the permanent archive, filed away and copied only when a cell divides. RNA became the

disposable working copy, transcribed fresh off the DNA archive whenever a particular protein is actually needed, then broken down and discarded. This division is the core of what molecular biologists later named the Central Dogma: information flows outward from a stable archive to disposable working copies, not the reverse.

The scale of that disposability has been measured directly, not merely assumed. A 2003 genome-wide study tracked the decay of every messenger RNA transcript in *Escherichia coli*, a workhorse gut bacterium, across 1,036 genes and 329 operons, an operon being a set of genes read off the DNA together as a single unit, and found the average message survives just 6.8 minutes before it is broken down. Compare that to the DNA archive itself, which persists across a cell's entire lifetime and is copied faithfully into every descendant. Two molecules, two wildly different lifespans, doing two very different jobs in the same cell.

That split is not treated as a weakness today. Cells now exploit short RNA lifespans on purpose, switching genes on and off quickly because the messenger evaporates fast once production stops. Whether that speed was the original reason short-lived RNA persisted, or whether it is an inherited holdover from RNA-world chemistry that later proved useful for an unrelated reason, remains an open question rather than a settled one.

The living fossil of the whole transition runs in the opposite direction from normal. Reverse transcriptase is an enzyme that reads RNA and writes DNA, the reverse of the usual archive-to-messenger flow, and it is not a curiosity confined to retroviruses like HIV. It turns up, repeatedly borrowed and repurposed, across every domain of life: in LINE-1 retrotransposons, mobile stretches of DNA that copy-paste themselves through an RNA go-between; in bacterial retrons and CRISPR-associated defense systems; in self-splicing group II introns; and, in eukaryotes with linear chromosomes, in telomerase, the enzyme that rebuilds the protective caps on chromosome ends.

A phylogenetic analysis posted in October 2025, still a preprint and not yet through formal peer review, traced telomerase's catalytic core to a

specific bacterial antiviral defense system named DRT10, a system that already used a noncoding RNA template to write short repeated stretches of DNA, sharing structural and sequence parallels in how that RNA template dictates the pattern of repeat addition, long before there were chromosome ends left needing protection. If that finding survives review, reverse transcription was not invented to solve the telomere problem at all; it was inherited from a far older bacterial toolkit and put to a new use. It is flagged here as provisional, since preprint evidence has not yet cleared the same bar as the rest of this leg's sources.

Every one of these clues points the same direction. DNA reads less like a replacement for RNA than a repair job built on top of it, still wearing RNA's fingerprints in its own raw materials and in enzymes that run the whole relationship backward. What DNA bought the lineage was time: an archive that could hold its shape for longer between copies. What it could not do on its own was read itself back out into a working protein. That reading job belonged to a separate piece of machinery, one already ancient by the time DNA took over the archive, and it is still running, essentially unchanged, in every cell alive.

Overtaken if: a natural cellular pathway were found that synthesizes DNA building blocks directly, without routing through an RNA precursor and ribonucleotide reductase, breaking the molecular fossil record tying DNA's own chemistry back to an RNA origin.

Sources

- **1993** · Lindahl, T., "Instability and decay of the primary structure of DNA" — [Nature 362:709–715, PMID 8469282](#)
- **2013** · Krokan, H.E. & Bjørås, M., "Base Excision Repair" — [Cold Spring Harbor Perspectives in Biology 5:a012583](#)
- **2015** · Lundin, D. et al., "The Origin and Evolution of Ribonucleotide Reduction" — [Life 5:604–636, PMC4390871](#)
- **2015** · Higgs, P.G. & Lehman, N., "The RNA World: molecular cooperation at the origins of life" — [Nature Reviews Genetics 16:7–17](#)
- **2003** · Selinger, D.W. et al., "Global RNA Half-Life Analysis in Escherichia coli Reveals Positional Patterns of Transcript Degradation" — [Genome Research 13:216–223](#)

- **2025** · (preprint) "Antiviral reverse transcriptases reveal the evolutionary origin of telomerase" — [bioRxiv](#) / [PMC12632851](#)

06. The Membrane and the Machine

Picture warm shallow water, or the plume of a hydrothermal vent, somewhere around four billion years ago. Fatty acid molecules, long carbon chains with a head that loves water and a tail that flees it, are crowding together in solution, and the moment enough of them gather, they curl spontaneously into hollow spheres. No gene directs this. No instruction exists yet to direct it. It is simply the shape this chemistry prefers once concentration and conditions allow it, and the hollow sphere it produces, called a vesicle, is the first rough draft of an inside separated from an outside.

Roy Black and Matthew Blosser's 2016 synthesis of the field, building on decades of work by pioneers such as Deamer, Luisi, Szostak and Rasmussen, describes these vesicles as more than passive bubbles. Built from single-chain fatty acids roughly ten to fifteen carbons long, they assemble across a wide range of conditions, a pH anywhere from about 6.5 to above 12, in water salted the way early oceans are thought to have been, with sodium, magnesium and calcium all present, and they form most readily around 70 degrees Celsius, scalding to the touch but unremarkable next to a hydrothermal source. Once formed, these vesicles actively bind and concentrate the raw building blocks of biology inside themselves: the individual letters that make up RNA, simple sugars, and, with somewhat less certainty, amino acids.

That concentrating effect matters more than it sounds. A self-copying molecule adrift in open water competes with nothing and protects nothing it makes; useful neighboring molecules simply drift away into the vastness of the ocean. Wrap the same chemistry inside a fatty membrane, and there is suddenly a boundary: an inside that can hoard its own building blocks, an outside that cannot immediately dilute them away. That boundary does real work for whatever chemistry it encloses, long before a single gene exists to build it on purpose. Every cell alive today still wears some descendant of that same boundary, a lipid bilayer: two layers of fatty molecules, tails pointing inward away from water, heads

pointing outward into it. Laboratory versions of these same simple vesicles today can still encapsulate and protect short strands of RNA, reproducing in a dish the proposed first step from four billion years ago.

How much a membrane like this could accomplish entirely unaided, before any genetically encoded help arrived in the form of membrane proteins the cell built for itself, is still argued over, and so is how strongly amino acids specifically stick to such an early membrane compared to how well nucleobases and sugars do. The first step's chemistry is settled and uncontested. How many further steps it could take alone is not.

A second puzzle concerns which kind of membrane came first. The two great domains of single-celled life today build their membranes from chemically mirror-image parts. Archaea use branched hydrocarbon chains, called isoprenoids, linked by a tough ether bond to one mirror-image form of a glycerol backbone. Bacteria, and Eukarya, which inherited their membrane chemistry from bacteria-like ancestors, use straight fatty acid chains linked by an easier-to-break ester bond to the opposite mirror-image form of that same backbone. This split was long assumed to have opened up only after LUCA, since no living organism was ever found straddling both chemistries at once.

In 2021, Laura Villanueva, Bas Dutilh and colleagues narrowed that assumption without closing it. They found a complete, apparently functional archaeal-style pathway for building ether lipids hiding inside the genomes of ordinary bacteria belonging to the FCB superphylum, a large bacterial grouping named for three of its member lineages, Fibrobacteres, Chlorobi and Bacteroidetes, with the relevant genes switched on in wild environmental samples, not merely present and silent. Moving those bacterial genes into laboratory *E. coli* produced cells running a genuinely mixed membrane, archaeal chemistry and bacterial chemistry side by side in the same cell wall. That does not prove LUCA itself had such a mixed, double-handed membrane; no living organism confirms that directly. It does show the two chemistries are far less walled off from each other than the textbook divide implied.

The machine that never changed

Whatever LUCA's own membrane looked like, something running inside it was already reading the frozen, sixty-four-word code from two legs back and using it to build proteins. That something is the ribosome: a large molecular machine, built mostly out of RNA with a supporting cast of proteins, that takes a strand of messenger RNA and, three letters at a time, strings together the matching chain of amino acids.

The proof that this machine is itself a working fossil of the RNA world arrived as a single, decisive structure. In 2000, Poul Nissen, Jeffrey Hansen, Nenad Ban, Peter Moore and Thomas Steitz solved the atomic structure of the ribosome's large subunit, resolved down to 2.4 angstroms, a unit small enough to distinguish individual atoms, using material from *Haloarcula marismortui*, a salt-loving archaeon. Their results appeared as two companion papers in the same issue of *Science*: one from Ban's team describing the subunit's overall architecture, the other from Nissen's team zeroing in on the active site itself. At the exact site where two amino acids are stitched together, called the peptidyl transferase center, no protein sits close enough to touch the reaction; the active site is built entirely from two RNA molecules, the 23S and 5S ribosomal RNAs. Thomas Cech, reviewing the result in a commentary in the same issue, summed it up in five words: "the ribosome is a ribozyme," a piece of RNA doing a job usually reserved for a protein enzyme.

The implication runs deep. RNA had to be capable of stitching amino acids together before proteins existed to do that job any other way. The ribosome's catalytic heart is not a metaphor for the RNA world; it is a still-working piece of it. That structural fact has not been seriously disputed since 2000. What remains under active study is exactly how such a ribozyme first assembled, and how it later acquired its surrounding shell of protein support.

That shell has its own history, and the ribosome's shape seems to remember building it. In 2009, Konstantin Bokov and Sergey Steinberg mapped the internal folding pattern of the large ribosomal RNA subunit and found it organized in concentric layers, an ancient core wrapped in

shell after shell added later, rather than as an arbitrary tangle of RNA. They concluded that the ribosome's origin "happened nearly four billion years ago," a figure reasoned out from structure rather than read off any rock. A later reanalysis, published in the *Journal of Molecular Evolution* under the title "Ancestral Insertions and Expansions of rRNA do not Support an Origin of the Ribosome in Its Peptidyl Transferase Center," pushed back on one detail, questioning whether that ancient core lines up exactly with the peptide-bond-forming site, without overturning the layered picture itself.

A separate 2015 study by Anton Petrov and colleagues built out a fuller six-phase growth sequence for the whole translation system: an ancestral RNA core first learned to fold correctly, then to catalyze a chemical reaction, then to pair up with a second subunit, then to co-evolve with its emerging protein partners, then to run on an energy-driven decoding cycle, and only last did today's proteinized outer surface take shape. Every ribosome alive still carries this same layered core; eukaryotic ribosomes, the kind found in cells with a nucleus, added two further and more recent layers of their own, long tentacle-like RNA extensions found nowhere else in the tree of life. Several such accretion models compete over the exact order of the layers, without disagreement on the broad direction: a small, ancient RNA core, growing outward into today's hybrid machine.

The reading step itself works through a division of labor between two kinds of molecule, as Miguel Angel Rubio Gomez and Michael Ibba's comprehensive 2020 review lays out. A family of enzymes called aminoacyl-tRNA synthetases, molecular matchmakers, one enzyme family shared across all of life, attach each of the twenty amino acids to its own correct transfer RNA, the adaptor molecule carrying the matching three-letter anticodon, checking and double-checking the pairing before release. Once that charging step is done correctly, the ribosome itself never rechecks the amino acid's chemical identity; it simply trusts the three-letter pairing between a codon on the messenger strand and an anticodon on the incoming transfer RNA, and forms the bond. Oddly, the matchmaker enzymes come in two structurally

unrelated families, Class I and Class II, ten enzymes each, that independently evolved to perform the identical job, a split whose reason is still an open question, one that reaches back toward the same pre-LUCA code-building era from two legs back.

Put the two halves of the reading system side by side, RNA-only catalysis at the ribosome's core, protein-based precision at the charging step, and the whole machine reads as a fossil record in its own right: an RNA-world veteran wearing tools built by the protein world it helped bring into being. That same machine, structurally recognizable down to its oldest core, runs unmodified in every domain of life today. Two independent inventions of a machine this specific, this old and this exact would be a coincidence nobody studying it has ever needed to invoke, which makes the ribosome, if anything, an even stronger argument for one shared ancestor than the frozen codon table alone.

A membrane to hold a self apart from the world. A frozen code to write proteins by. A machine old enough to remember reading the first of them. By roughly four billion years ago, give or take the usual margins, the essential parts were all in place. What was not yet in place was scale, something to make energy abundant enough to build a far larger cell out of these same old parts. That arrives next, when one cell does something stranger than inventing any new molecule: it swallows another cell whole, and keeps it.

Overtured if: a ribosome's peptidyl transferase center were found built from protein rather than RNA in any lineage, breaking the claim that peptide-bond catalysis is a universal, inherited relic of the RNA world.

Sources

- **2016** · Black, R.A. & Blosser, M.C., "A Self-Assembled Aggregate Composed of a Fatty Acid Membrane and the Building Blocks of Biological Polymers Provides a First Step in the Emergence of protocells" — [Life 6:33](#)

- **2021** · Villanueva, L. et al., "Bridging the membrane lipid divide: bacteria of the FCB group superphylum have the potential to synthesize archaeal ether lipids" — [The ISME Journal 15:168–182](#)
- **2000** · Nissen, P., Hansen, J., Ban, N., Moore, P.B., Steitz, T.A., "The structural basis of ribosome activity in peptide bond synthesis" — [Science 289:920–930](#), PMID 10937990
- **2009** · Bokov, K. & Steinberg, S.V., "A hierarchical model for evolution of 23S ribosomal RNA" — [Nature 457:977–980](#)
- **2015** · Petrov, A.S. et al., "History of the Ribosome and the Origin of Translation" — [PNAS 112:15396–15401](#), PMID 26621738
- **2020** · Rubio Gomez, M.A. & Ibba, M., "Aminoacyl-tRNA synthetases" — [RNA 26:910–936](#), PMC7373986

07. The Great Merger

About 2.2 billion years ago, in an ocean lit by a fainter sun, under a sky still adjusting to a new supply of oxygen, something merged that is still running today, inside a beating human heart. The era is the Paleoproterozoic, the opening chapter of Earth's Proterozoic history — a stretch of time roughly two billion years after LUCA, life's last universal common ancestor, first appeared. Nothing with eyes existed yet. Nothing had a nerve. What happened in that ocean was smaller than one cell noticing another, and it turned out to be larger than almost anything that came after it.

No fossil records the moment itself; cells that small and that old almost never leave a shape behind. Two rival reconstructions compete for what the scene actually looked like. In one, a larger host cell engulfed a smaller bacterium whole, through a process called phagocytosis — cell-eating — and then, instead of digesting its catch, kept it alive inside its own body, permanently. In the other, called the hydrogen hypothesis, there was no ambush at all: an archaeon (one of life's three great domains, alongside bacteria and the more complex cells that would follow) that could only survive by consuming hydrogen gas settled into permanent partnership beside a bacterium that released hydrogen as a waste product of its own breathing, until the two became so mutually dependent that separation stopped being an option. William F. Martin and Miklós Müller proposed the hydrogen version in 1998, arguing that energetic interdependence, not predation, was the glue from the very start. It still competes today with the classical predation story and with a third idea — that the partnership began as protection from oxygen poisoning — rather than having settled into one accepted account.

Whichever scene is closer to true, the general idea has a name and a modern author. Endosymbiosis means, literally, living together, one cell inside another, permanently. The serial endosymbiotic theory — proposing that mitochondria (the compartments that generate a cell's usable energy), plastids (the compartments that carry out photosynthesis

in plant and algal cells), and, more speculatively, the whip-like flagellum all began as free-living bacteria taken up and retained by a host cell — was assembled into scientific respectability by the biologist Lynn Margulis, publishing in 1967 under her name at the time, Lynn Sagan. She built on ideas about symbiosis inside cells first raised in the early twentieth century by Mereschkowski and Wallin. Her paper carried no new experiment — it was a synthesis, a case built from existing evidence pointed in a new direction — and more than a dozen journals rejected it before the *Journal of Theoretical Biology* finally printed it.

Time has sorted her claims into two very different piles. The mitochondrial and plastid parts of the theory are now mainstream science, accepted across cell biology, evidenced by every eukaryotic cell (a cell with a nucleus and internal compartments, the design shared by everything from yeast to whales) that still carries mitochondria, and by every photosynthetic eukaryote that still carries a plastid descended from a captured cyanobacterium. The flagellum-from-spirochete part — the idea that the whip-like tail some cells use to swim also began as a captured bacterium, in this case a corkscrew-shaped spirochete — never gathered comparable support, because no one has found the genomic evidence for it that turned up so clearly for mitochondria. A theory, it turns out, can be half right and still remake a field.

Margulis dated none of this; her argument was about mechanism, not timing. The dates arrived decades later, from a study led by Paschalia Kapli, Davide Pisani and Philip C. J. Donoghue, of the Bristol Palaeobiology Group, published in *Nature* with an online release in December 2025 and print in 2026. Their method, a molecular clock, treats mutations as a kind of metronome: genes accumulate changes at roughly steady rates, so comparing how different two copies of a gene have become can estimate how long ago they split from a shared original. The team built a hundred and thirty-five time-resolved family trees for individual genes — ninety-five rooted in archaea, forty rooted in bacteria — layered across a master tree of sixty-two marker genes, and calibrated the whole structure against eighteen independent fossils.

What emerged reorders the story. The host lineage, which the paper labels nFECA – the nuclear first eukaryotic common ancestor, in its own terminology – split away from the rest of the Asgard archaea, its closest surviving relatives, somewhere between 3.05 and 2.79 billion years ago. Critically, this host had already built itself an internal skeleton (a cytoskeleton), internal membrane compartments (an endomembrane system), and the cellular engulfing machinery of phagocytosis, before it ever acquired a mitochondrion. Only later did the mitochondrial ancestor – mFECA, the mitochondrial first eukaryotic common ancestor – split from its own closest relatives, a large bacterial family called the Alphaproteobacteria, sometime between 2.37 and 2.13 billion years ago. Specific gene duplications tied to importing proteins into the mitochondrion, in genes named TIM14 and TIM44 – essentially molecular doorkeepers that let proteins cross the mitochondrion's inner membrane – push the working estimate for the actual onset of the merger to around 2.2 billion years ago. Only after the merger was fully absorbed into a single working cell did the last eukaryotic common ancestor, LECA, radiate into the lineages that would become every plant, animal, fungus and protist alive today, between 1.80 and 1.67 billion years ago.

Living relatives of both halves of the merger still exist. Asgard archaea – a group that includes the aptly named Lokiarchaeota and Heimdallarchaeia – are the closest surviving kin of the ancient host cell. The Alphaproteobacteria remain the closest surviving kin of the ancient mitochondrial symbiont. The order this study reconstructs, a complex host first and a captured bacterium second, directly challenges older mitochondrion-early scenarios that pictured a simple archaeon transformed into complexity only after picking up its symbiont; the 2026 paper explicitly argues against that older picture, though the precise timeline, and the deeper root of the archaeal-bacterial tree itself, remain live questions, discussed in an accompanying commentary by Hua in the same journal family.

The Energy Behind Complexity

None of this timing explains why the merger mattered so much. For that, the clearest argument comes from Nick Lane and William F. Martin, of University College London and Heinrich Heine University Düsseldorf, in a 2010 paper in *Nature*. Their case is bioenergetic, built from a mismatch of geometry. A cell makes its usable energy currency, a molecule called ATP, across membrane surfaces, but the cost of building and running a larger genome scales with the volume of the whole cell. Surface area and volume do not grow at the same rate — area lags behind as an object gets bigger — so a bacterium scaled up to the size of a typical complex cell would starve its own genes of power long before it got there. By multiplying the amount of energy-generating membrane packed inside a single cell, without also multiplying that cell's outer surface, mitochondrial endosymbiosis broke through that ceiling. Lane and Martin calculate the resulting gain at roughly two hundred thousand-fold in the number of genes a cell could afford to keep switched on.

This is not settled science. Michael Lynch and Georgi K. Marinov called the entire energy-per-gene framing ecologically and evolutionarily meaningless, arguing that once cell volume is properly accounted for, mitochondria do not actually expand a cell's bioenergetic capacity at all. Lane and Martin answered in 2016, in a PNAS paper whose title turned the criticism back around, framing mitochondrial complexity as a kind of evolutionary deficit spending; Lynch and Marinov published a counter-reply the same year. The dispute remains open and active — a reminder that a mechanism can be repeated confidently in popular science while professional disagreement about it continues underneath.

The clearest surviving proof that any of this happened at all is not an argument but an object: mitochondrial DNA. It is small, circular, and shaped like a bacterial chromosome rather than like the long chromosomes coiled inside a cell's nucleus, because it is, in every meaningful sense, a shrunken bacterial chromosome that never finished being absorbed. In humans it runs to 16,569 base pairs (the paired chemical letters, A-T and G-C, that spell out DNA's instructions) and

carries just 37 genes: 13 for proteins used in respiration, 22 for the small RNA adaptors called tRNAs, and 2 for the RNA scaffolding of the ribosome. Over the two billion years since the merger, most of the original bacterial genome moved out of the mitochondrion and into the cell's own nucleus; what is left behind is itself a kind of clock, its smallness a rough gauge of how much time has passed since the takeover began.

Even the question of exactly which bacteria mitochondria are closest to remains disputed, three ways at once. Rickettsiales – the bacterial family that includes the microbes responsible for typhus and spotted fever – has long been the classical candidate. A 2022 reanalysis by Sergio A. Muñoz-Gómez and colleagues, comparing 108 mitochondrial proteins with modeling designed to correct for compositional differences between distantly related genomes, instead places mitochondria as sister to the Alphaproteobacteria as a whole, not to Rickettsiales specifically. Joran Martijn and colleagues, in 2018, argued for something stranger still: that the true sister lineage split off even earlier, and is not represented by any alphaproteobacterium currently sampled and sequenced at all – a relative that has, so far, evaded discovery. All three positions remain in circulation, and the disagreement turns on how each team samples its bacteria and models the statistical noise of two billion years.

What is certain is simpler than any of the three positions. The merger happened; energy became abundant inside a single cell for the first time; and every large, complex cell alive today descends from that one partnership, however it began. Nothing that followed – no thought, no sensation, no signal passed from one cell to another on purpose – was yet possible in that Paleoproterozoic ocean. But the energy to build such things was, for the first time, sitting inside a cell with room to spend it. The molecules that would one day carry feeling had already begun appearing in cells simpler than this one, and it is to them, still blind, still alone, that the story turns next.

Overtaken if: a robust phylogenomic dataset reverses the sequence the 2026 Donoghue, Kapli and Pisani gene-duplication study reports –

showing the mitochondrial ancestor's arrival preceding, rather than following, the evolution of the host lineage's cytoskeleton, endomembrane system and phagocytotic machinery.

Sources

- **1967** · Lynn Sagan, "On the Origin of Mitosing Cells" — [Journal of Theoretical Biology](#), via [Embryo Project Encyclopedia](#); 50-years retrospective — [Molecular Biology of the Cell](#)
- **2026** · Donoghue, Kapli, Pisani et al., "Dated gene duplications elucidate the evolutionary assembly of eukaryotes" — [Nature 650](#), 129–140; commentary — [Hua](#), [Nature Genetics 58](#), 13 (2026)
- **2010** · Lane & Martin, "The energetics of genome complexity" — [Nature 467](#), 929–934; rebuttal — [Lynch & Marinov](#), [PNAS 113](#)(6):E668–E669 (2016)
- **1998** · Martin & Müller, "The hydrogen hypothesis for the first eukaryote" — [Nature 392](#), 37–41
- **2022** · Muñoz-Gómez et al., "Site-and-branch-heterogeneous analyses of an expanded dataset favour mitochondria as sister to known Alphaproteobacteria" — [Nature Ecology & Evolution 6](#), 253–262; mtDNA structure — [Anderson et al.](#), [PMC5267354](#)
- **2018** · Martijn et al., "Deep mitochondrial origin outside the sampled alphaproteobacteria" — [Nature 557](#), 101–105

08. Messengers Before Minds

On a sun-warmed tidal flat, somewhere between three billion and two and a half billion years ago, as the Archean eon was giving way to the Proterozoic, a mat of cyanobacteria — the blue-green, oxygen-releasing bacteria that were quietly remaking Earth's atmosphere — is manufacturing its own poison. Photosynthesis, the chemical trick of turning sunlight into sugar, throws off free radicals as a byproduct: unstable molecule fragments that damage whatever they touch. Every cell in the mat bathes in its own exhaust. Inside those cells, a small molecule is already on duty, neutralizing the damage before it spreads. Its name is melatonin. Not one nerve cell exists anywhere on the planet yet, and nothing that could be called a mind is remotely close, yet a molecule later associated, in animals, with sleep and mood is already doing exactly the job chemistry demands of it: mopping up radicals in a cell too simple to dream.

Serotonin, dopamine, histamine and melatonin are usually introduced as neurotransmitters — chemical messengers a nerve cell releases to signal its neighbor across a tiny gap called a synapse. That description is accurate for what these molecules do in animals with nervous systems. On the scale of deep time, it is almost an afterthought. Each of these molecules did other jobs first — metabolic, defensive, regulatory — for a very long stretch of Earth's history before any nervous system existed to recruit them for passing signals between neurons. The chemistry of feeling is older than feeling itself.

The melatonin case is argued in the most detail by Russel J. Reiter, Dun Xian Tan, Lucien C. Manchester and colleagues at UT Health San Antonio, synthesizing comparative biochemistry across 2015 and 2016. Melatonin's molecular structure, they note, has stayed essentially unchanged from cyanobacteria all the way to humans — a chemical fossil hiding in plain sight inside every cell that still makes it. Their proposed date for its origin, roughly 3.0 to 2.5 billion years ago, is pinned to the Great Oxidation Event of about 2.5 to 2.0 billion years ago, when rising

photosynthetic oxygen would have first created strong selection pressure for exactly this kind of antioxidant. That date is a geological inference, not a molecular-clock measurement, and should be read as a plausible order of magnitude rather than a calibrated year. Cyanobacteria alive today still manufacture melatonin non-neurally, for the same antioxidant purpose their ancestors needed it for; the pathway is proposed to have entered eukaryotic cells by the same cyanobacterial endosymbiosis that, elsewhere in the cell, produced the plastid.

Serotonin and dopamine share a deeper ancestry, traced through the enzymes that build them rather than through the molecules themselves. Tryptophan hydroxylase and tyrosine hydroxylase – the two enzymes that perform the rate-limiting step in making serotonin and dopamine – belong, together with a third enzyme called phenylalanine hydroxylase, to a single family of pterin-dependent hydroxylases (pterin being a small helper molecule the enzyme needs to do its chemistry). Linda Skjærven, Knut Teigen and Aurora Martinez, at the University of Bergen, synthesized the case in 2014, building on chromosome-mapping work from 1987 that had located the human genes for tryptophan hydroxylase and tyrosine hydroxylase together on chromosome 11, separate from phenylalanine hydroxylase on chromosome 12. Phenylalanine hydroxylase is the ancestral member of the family; the whole animal lineage of these enzymes descends, through at least two gene duplications and one chromosomal relocation, from a hydroxylase family already present in prokaryotes – cells with no nucleus, the simple design shared by every bacterium – long before animals existed to need serotonin or dopamine at all. *Dictyostelium discoideum*, a social amoeba, carries an independently studied version of this hydroxylase, still used today to help root the family tree.

Whether bacteria actually make serotonin with this inherited machinery, rather than merely carrying look-alike genes, is unresolved. A 2022 review led by Luqman, Blackshaw and Wei-Cheng Liu surveyed bacterial genomes and found the enzymes capable of making serotonin – aromatic amino acid hydroxylases and decarboxylases – widely present, while confirmed serotonin production itself was absent in most of the species

carrying those genes. The honest reading, which the review states plainly, is a gap rather than an answer: bacteria may be reaching the same molecule by a different, non-homologous chemical route, rather than running an inherited animal-style pathway. What is not in doubt is that specific gut bacteria alive today, including spore-forming species and Akkermansia-type microbes, measurably influence how much serotonin their animal hosts produce – a functional link at least as old as the animal gut itself.

Plants make dopamine too, and by two different routes that never both run at once in the same species. One path hydroxylates tyrosine into L-DOPA before decarboxylating it into dopamine, the same order animals use. The other decarboxylates tyrosine into tyramine first, then hydroxylates it into dopamine using enzymes called polyphenol oxidases. Which route a given plant uses is species-specific: a 2025 study found the model grass *Brachypodium distachyon* relies on the tyramine-first route as its primary pathway. Roshchina and colleagues, reviewing the wider field in 2020, catalogue dopamine's plant jobs as growth and development signaling, antioxidant defense, and accumulation under both biotic stress (attack from other organisms) and abiotic stress (heat, drought, salt). Banana, potato, various legumes and *Brachypodium* are named, characterized examples. No molecular clock has dated plant dopamine biosynthesis specifically; the shared ancestry of its enzymes with the animal versions implies an origin before the plant and animal lineages split, but that is an inference from enzyme-family relationships, not a measured date.

A Patchier Inheritance

Histamine breaks the pattern the other three molecules suggest, and the break is worth naming carefully. The amino acid it is built from, histidine, has an origin as deep as origins get: the pathway that builds histidine from scratch is inferred, from comparative genomics of the *his* operon across Archaea, Bacteria and Eukarya by Marco Fondi, Renato Fani and colleagues in 2009, to have been fully assembled before LUCA, the last

universal common ancestor of all life, roughly 3.5 to 4 billion years ago, in the Archean eon or earlier, because essentially the same pathway shows up across all three domains of life. Histamine itself is a different story. The enzyme that removes a carbon-dioxide group from histidine to make histamine, called histidine decarboxylase, exists as two structurally unrelated enzyme families that evolved independently of one another – one dependent on a helper molecule called pyridoxal-5'-phosphate, found mainly in Gram-negative bacteria, the other a self-processing, pyruvoyl-dependent type found mainly in Gram-positive bacteria such as lactic acid bacteria.

The clearest evidence against a deep, uniform origin for histamine comes from a 2017 study by Daniel Wüthrich, H  l  ne Berthoud, Stefan Irmler and colleagues at Agroscope and the University of Bern, comparing the histamine-producing gene cluster across strains of *Lactobacillus parabuchneri*. The genes showed close to ninety-nine percent nucleotide identity to homologs in other *Lactobacillus* and in *Oenococcus* – a similarity so high it signals recent horizontal gene transfer, genes moving sideways between species, rather than billions of years of vertical, parent-to-offspring inheritance. The lesson generalizes: histamine-producing capacity is patchy and mobile across the bacterial tree, not evenly and anciently distributed the way the histidine pathway beneath it is. It stands as a caution against assuming every biogenic amine is as old as its parent amino acid. *Lactobacillus parabuchneri*, *Tetragenococcus halophilus* and *Oenococcus oeni*, all found today in fermented foods, are the clearest living non-neural histamine producers. Plants accumulate histamine under stress but appear to lack an active histidine-decarboxylase pathway of their own; a candidate gene in *Arabidopsis*, once proposed as a plant histidine decarboxylase, was later shown to actually be a serine decarboxylase, with no activity on histidine at all.

The single clearest proof that this entire chemical kit predates nervous systems comes from a single-celled organism with no nervous system to speak of. Gyula Csaba, at Semmelweis University, spent decades studying the ciliate *Tetrahymena* — a single cell covered in hair-like cilia, common in freshwater ponds — and reviewed the results in 2015. *Tetrahymena*

synthesizes, stores and secretes nearly the entire animal signaling kit at once: histamine, serotonin, epinephrine, dopamine and melatonin, together with binding sites resembling the G-protein-coupled receptors that animal nervous systems use to detect these same molecules. It uses this chemistry to regulate its own phagocytosis (engulfing food), cell division, the regrowth of its cilia, glucose metabolism, and chemotaxis – swimming toward or away from a chemical cue. All of this happens, Csaba's review argues, in an organism separated from any nervous system by hundreds of millions of years and by an entire kingdom of life. *Tetrahymena pyriformis* and its relatives, in the group called ciliates, are the living species in which this has been directly measured; whether their amine receptors are truly the same ancestral proteins later used by animal nerves, or a separate, convergent invention, remains an open question no single study has settled.

None of these molecules, across three and a half billion years of bacteria, plants and single cells, ever produced a feeling. A cyanobacterium mopping up radicals with melatonin is not calm. A fermenting *Lactobacillus* making histamine is not alarmed. *Tetrahymena* swimming up a gradient of its own dopamine is not wanting anything, in any sense that would be recognized inside a body with a brain. What existed, everywhere these molecules turn up, was chemistry doing metabolic and regulatory work with no experience attached to it. The experience would need something none of these organisms had: cells specialized for nothing but carrying a signal, wired together into a network. That network, and the animals built around it, is what appears next.

Overtured if: a bacterial species is confirmed to synthesize serotonin through the same homologous aromatic-amino-acid-hydroxylase and decarboxylase pathway animals use, rather than a convergent, non-homologous route, closing the evidence gap the 2022 Luqman, Blackshaw and Liu review leaves open.

Sources

- **2022** · "Enzyme Promiscuity in Serotonin Biosynthesis, From Bacteria to Plants and Humans" — [Frontiers in Microbiology 13:873555 / PMC9048412](#)
- **2015** · Manchester et al., "Melatonin: an ancient molecule that makes oxygen metabolically tolerable" — [Journal of Pineal Research 59\(4\), 403–419](#); mechanism review — [Reiter et al., Cellular and Molecular Life Sciences 74, 3863–3881 \(2017\)](#)
- **2020** · "Functions of dopamine in plants: a review" — [Plant Signaling & Behavior](#); recent mechanistic study — [bioRxiv, "Dopamine biosynthesis in plants" \(2025\)](#)
- **2017** · Wüthrich, Berthoud, Wechsler, Eugster, Irmeler & Bruggmann, "The Histidine Decarboxylase Gene Cluster of *Lactobacillus parabuchneri* Was Gained by Horizontal Gene Transfer" — [Frontiers in Microbiology 8:218](#); LUCA-era histidine pathway — [Fondi, Emiliani, Liò, Gribaldo & Fani, Journal of Molecular Evolution 69, 512–526 \(2009\)](#)
- **2015** · Csaba, "Biogenic amines at a low level of evolution: production, functions and regulation in the unicellular *Tetrahymena*" — [Acta Microbiologica et Immunologica Hungarica 62\(2\), 93–108](#)
- **2014** · Skjærven, Teigen & Martinez, "Structure–Function Relationships in the Aromatic Amino Acid Hydroxylases Enzyme Family: Evolutionary Insights" — [Wiley Online Library / eLS](#); locus history — [Craig et al., PMID 2889273 \(1987\)](#)

09. The First Animals

In the Mackenzie Mountains of northwestern Canada, a band of ancient reef limestone called the Little Dal Group holds branching, fibrous networks entombed inside 890-million-year-old microbial rock — the Tonian period, the opening act of the Neoproterozoic Era. Elizabeth C. Turner, of Laurentian University, published an analysis of these structures in *Nature* in 2021, interpreting them as fossilized tissue from a demosponge, one of the simplest kinds of sponge. If the interpretation holds, these would be animal fossils roughly 350 million years older than the next-oldest undisputed sponge body fossils, and they would predate the rise of atmospheric oxygen long thought necessary for animal life by about 90 million years. Turner herself, and outside commentators, are careful to call the finding tentative and dependent on later verification: no unambiguous spicules, the mineral or protein skeletal needles that support a sponge's body, were found alongside the fibrous networks, and the structures could instead be abiotic mineral growths or ordinary microbial-mat artifacts. Best treated as a disputed maximum, not an accepted date.

Firmer footing comes from chemistry rather than shape. A steroid molecule called 26-methylstigmastane is made, among living organisms, only by demosponges; its presence in ancient rock, reported by Love and colleagues in 2009, places definite demosponges at least 635 million years old. Demospongiae remains the most species-rich sponge class alive today, and one of its living members, *Amphimedon queenslandica*, a sponge native to the Great Barrier Reef, has become the standard genomic model for the whole basal sponge lineage — the organism whose DNA keeps answering questions no fossil can.

For a date on the whole animal kingdom rather than sponges alone, the reference point is a 2015 study by Mario dos Reis, Yuttapong Thawornwattana, Konstantinos Angelis, Maximilian J. Telford, Philip C. J. Donoghue and Ziheng Yang. Their method, a Bayesian molecular clock, treats the accumulated differences between living species' genes as a

kind of stopwatch, converted into calendar time by anchoring select branches of the tree to fossils of known age. Using 203 nuclear proteins across 71 species and 38,577 amino-acid sites, and testing four different fossil-calibration strategies against one another, they placed the origin of crown Metazoa — the group containing every animal alive today plus its most recent common ancestor — in a 95 percent credibility interval of 649.8 to 833.5 million years ago, spanning the boundary between the Tonian and Cryogenian periods, the era of Snowball Earth glaciations.

The same study dated crown Eumetazoa (cnidarians plus bilaterians, effectively every animal with true tissue layers) to 625.9–746.2 million years ago, crown Cnidaria (corals, jellyfish, sea anemones) to 531.5–641.8 million years ago, and crown Bilateria (everything else, including every vertebrate) to 595.7–688.3 million years ago. The paper's real headline, though, is a warning rather than a number: its own authors conclude that building firm evolutionary narratives on molecular-clock timescales this deep remains premature, because unavoidable violations of clock-like behavior over such vast stretches of time, combined with genuine uncertainty in fossil calibration, keep the true dates from narrowing further. Before this sensitivity analysis, competing estimates for crown Metazoa in the literature had ranged as widely as 1,298 million years to 615 million years — a spread the dos Reis team's wide, honest ranges were built specifically to explain.

Among the animals whose age these methods try to pin down is one of the strangest: *Trichoplax adhaerens*, representative of the group called Placozoa. It is flat, about a millimeter across, brainless, gutless and muscle-less, built from only a handful of cell types — among the simplest known free-living animals. Yet its genome, sequenced and described in 2008 by Mansi Srivastava, Emina Begovic, Jarrod Chapman, Daniel S. Rokhsar and colleagues at the DOE Joint Genome Institute and UC Berkeley, runs to roughly 98 million base pairs and encodes a surprisingly rich set of transcription factors and signaling-pathway genes, the kind otherwise associated with the complex development of animals with true tissues. Molecular phylogenetics places the placozoan lineage's divergence from other animals at more than 600 million years ago; the

genome paper's own analysis placed Placozoa inside a broader animal clade alongside cnidarians and bilaterians, with sponges branching off earliest of all. Several additional, genetically distinct placozoan lineages, hiding under what had looked like one species, have since been identified in comparative genomic work published in 2018. Exactly where Placozoa sits relative to sponges, ctenophores and the rest remains one of the least settled nodes in the whole animal tree.

A demosponge without a single neuron offers the clearest hint of what came before nervous systems existed. *Amphimedon queenslandica*, the same sponge used as the genomic model above, was shown in 2007 by Onur Sakarya, Kenneth W. Armstrong, Kevin J. Peterson and colleagues to carry, in its genome, a nearly complete postsynaptic protein scaffold — the molecular anchoring structure that, in an animal with a nervous system, holds signal-receiving proteins in place on the receiving side of a synapse. This includes the scaffolding gene *dlg*, known in humans as PSD-95, with the exact residues that contact its binding partner 100 percent identical between sponge and human, and co-expressed together during both larval and adult stages. Missing entirely from the sponge genome, however, are the actual ionotropic glutamate receptors that would dock onto that scaffold inside a real neuron. The reading is that the molecular hardware for a synapse assembled before neurons themselves existed, with a key receptor component added later, elsewhere in the tree. Follow-up work in 2019 found the same genes co-expressed in adult sponge choanocytes, the hair-ringed cells sponges use to filter food from water, reinforcing a pre-neural, non-neuronal function.

Two Candidates for First

Ctenophores, the comb jellies, complicate the picture further. Leonid L. Moroz, of the University of Florida, and independently Joseph F. Ryan and colleagues, sequenced ctenophore genomes — *Pleurobrachia bachei* by Moroz's team, and *Mnemiopsis leidyi*, at roughly 150 million bases, by Ryan's — both published in 2013 and 2014. Ctenophores plainly have a

working nervous system. But it lacks every classic neurotransmitter found in other animals with nerves: no acetylcholine, serotonin, histamine, dopamine, adrenaline, noradrenaline or octopamine signaling turned up in either genome. In their place, an unprecedented diversification of ionotropic glutamate receptors – the same receptor family conspicuously missing from the sponge genome – appears to carry the signaling load instead. The implication argued by both teams is that ctenophore nervous systems were built from a different molecular parts list entirely, meaning they may have arisen independently of, rather than by inheritance from, the nervous systems of every other animal.

Whether that independence is real depends on a question still unresolved: which lineage, sponges or ctenophores, branched off earliest from every other living animal. The answer decides whether nervous systems evolved exactly once, later lost in sponges and placozoans, or at least twice, independently, once in ctenophores and again in the shared ancestor of cnidarians and bilaterians. Roberto Feuda, Davide Pisani and colleagues argue for sponges as the earliest branch, having reanalyzed multi-gene phylogenomic matrices in 2017 using improved models of amino-acid compositional differences between species, and concluding that the rival ctenophore-first result is a "tree-reconstruction artifact" produced by inadequate modeling. Joseph F. Ryan, Leonid Moroz, Casey Dunn and colleagues argue the opposite. In 2023, Darrin T. Schultz, Stephen A. Smith, Ryan and colleagues mustered a different kind of evidence for the ctenophore-first position: patterns of conserved gene neighborhood, called synteny, in which particular genes stay next to the same neighboring genes across enormous stretches of evolutionary time. Sponges, cnidarians and bilaterians share certain rare synteny patterns that ctenophores lack, a marker the Schultz team argues resists the modeling problems that plague amino-acid-based methods. A 2023 exchange in *Nature Communications* shows neither camp treats the question as closed.

The fossil record, when it finally arrives for ctenophores, arrives much later than any of these molecular estimates. The oldest confirmed ctenophore body fossil is an embryo from the Meishucun assemblage in

Shaanxi, China, confirmed as organic rather than purely mineral preservation by Raman spectroscopy in 2007, and dated to the Early Cambrian, roughly 540 million years ago. Adult forms from the Chengjiang biota in Yunnan, China — among them *Maotianoascus octonarius*, built with skeletal support and as many as 24 comb rows for swimming, compared with just 8 in living species — date to about 518 million years ago. Both dates fall more than 100 million years after the Cryogenian-period origin the molecular clock implies for crown Metazoa, a textbook illustration of the gap that regularly opens between when a lineage's genes say it must already have existed and when the fossil record actually catches it in stone.

That fossil record, once it starts, is thin and easy to misread. A separate Chengjiang fossil, *Stromatoveris psygmoglena*, once proposed as a ctenophore relative, is now argued by some researchers to belong instead to an unrelated, frond-shaped group left over from the earlier Ediacaran period, called the *Petalonamae* — a reminder of how easily soft, frond-like Cambrian fossils get miscategorized. A further stem-ctenophore fossil, roughly 358 million years old, found at Miguasha in Canada and dated to the Late Devonian, shows the lineage persisting, sparsely documented, across the quarter-billion years between the Cambrian and the more familiar comb jellies alive today: some 100 to 150 living species, among them *Mnemiopsis leidyi* and *Pleurobrachia bachei*, all still running on the glutamatergic signaling pattern their Cambrian-and-older ancestors apparently invented.

None of these animals, brainless sponge, cell-poor placozoan, or nerve-net-bearing ctenophore, had anything resembling a mind. What they had, whichever of them came first, was ancient signaling chemistry, already billions of years old, finally being wired for the first time into cells whose only job was passing a signal on. The molecules that follow next are not simple amines mopping up radicals or nudging a cell's metabolism. They are peptides, small chains of amino acids, built for a purpose none of these early animals yet needed: holding one animal to another.

Overtured if: a phylogenomic method accepted by both the sponge-first camp of Feuda and Pisani and the ctenophore-first camp of Ryan, Moroz and Schultz converges on a single, mutually agreed sister lineage to all other animals, settling whether nervous systems arose once or independently at least twice.

Sources

- **2021** · Turner, "Possible poriferan body fossils in early Neoproterozoic microbial reefs" — [Nature 596, 87–91](#); biomarker evidence — [Love et al., Nature 457, 718–721 \(2009\)](#)
- **2015** · dos Reis, Thawornwattana, Angelis, Telford, Donoghue & Yang, "Uncertainty in the Timing of Origin of Animals and the Limits of Precision in Molecular Timescales" — [Current Biology 25\(22\), 2939–2950 / PMC4651906](#)
- **2008** · Srivastava, Begovic, Chapman, Putnam, Hellsten, Kawashima, Rokhsar et al., "The Trichoplax genome and the nature of placozoans" — [Nature 454, 955–960](#)
- **2007** · Sakarya et al., "A Post-Synaptic Scaffold at the Origin of the Animal Kingdom" — [PLOS ONE 2\(6\), e506 / PMC1876816](#); follow-up — [Najle et al., Scientific Reports \(2019\)](#)
- **2014** · Moroz et al., "The ctenophore genome and the evolutionary origins of neural systems" — [Nature 510, 109–114](#); **2013** · Ryan et al., "The genome of the ctenophore Mnemiopsis leidyi" — [Science 342\(6164\), 1242592](#)
- **2017** · Feuda et al., "Improved Modeling of Compositional Heterogeneity Supports Sponges as Sister to All Other Animals" — [Current Biology 27\(24\), 3864–3870](#); countered by — [Schultz, Haddock, Bredeson, Green, Simakov & Rokhsar, "Ancient gene linkages support ctenophores as sister to other animals," Nature 618, 110–117 \(2023\)](#)

10. The Bonding Peptides

On a warm night beside a rain pool in Panama, a male túngara frog inflates his throat and calls. The call is the whole of his advertisement: a whine sliding into a chuck, aimed at any female within earshot. Researchers who inject him with a single molecule called vasotocin do not change his desire to call. They change the call itself — its fine structure, and how readily other males move toward it. A single messenger molecule, hardly more than a handful of atoms strung into a chain of nine amino acids, is enough to retune a piece of mating behavior that evolution spent hundreds of millions of years building.

Nine amino acids joined end to end is called a nonapeptide — a name that describes only the length of the chain, not its job. In mammals, two members of this short family carry familiar names: oxytocin, tied to bonding, birth, and milk letdown, and vasopressin, tied to holding onto water in the kidney and to some of the same social wiring. Popular science writing sometimes calls oxytocin the love hormone, which oversells a molecule that also does plumbing work in the bladder and blood vessels. The underlying chemistry, though, is not oversold at all. The same nine-link peptide, wearing different names, still runs through worms, snails, fish, frogs, and birds — a thread of attachment chemistry older than bone.

The mammalian pair did not appear from nothing. Daniel Ocampo Daza, Christina Bergqvist, and Dan Larhammar, working across Uppsala University and UC Merced, traced the six-gene family of oxytocin and vasotocin receptors found in living jawed vertebrates back to one ancestral gene pair, sitting side by side on one ancient chromosome. That pair became six through two separate rounds of whole-genome duplication — an event in which an organism's entire genetic library is copied all at once, not gene by gene, doubling every chromosome in a single generation. Most duplicate genes made this way are later lost, deleted as redundant; a few survive by finding new work. The oxytocin/

vasotocin receptor family is a case where several duplicates survived, each specializing.

The evidence for two duplications, not one, is written into the chromosomes themselves as synteny — the tendency of genes to keep the same neighbors and the same order along a chromosome even after millions of years of separate evolution. The surviving genes still sit in matching arrangements on four different chromosomes, a fourfold pattern that is very hard to produce by any process except two doublings of the whole set. The study appeared online on 22 November 2021 and in print on 3 February 2022, in *Frontiers in Endocrinology*, volume 12, article 792644, drawing on freshly sequenced genomes — the reedfish *Erpetoichthys calabaricus*, the thorny skate *Amblyraja radiata*, the Western clawed frog *Xenopus tropicalis* — to confirm the fourfold symmetry across a wider slice of the vertebrate tree than earlier work had reached.

On the question of age, the authors are candid about the limits of their own method. A naive molecular clock — reading the number of accumulated mutations as if it were a stopwatch — applied to this gene family, in their own assessment, is often "overestimated," sometimes reaching back beyond the protostome-deuterostome split: the boundary where one great branch of the animal tree, the protostomes (worms, snails, insects), separated from the other, the deuterostomes (sea stars, fish, and eventually us), conventionally placed just over 600 million years ago. Rather than trust that clock on its own, the team bounded the duplications between two independently dated events instead — no older than the split separating vertebrates from tunicates, no younger than the point where bony and cartilaginous fish parted ways. That places the ancestral nonapeptide gene system somewhere between 600 and 450 million years old: a wide window, honestly reported as a window rather than rounded to one reassuring number.

Not every specialist agrees on how it happened. A rival reconstruction, built by Theofanopoulou and colleagues from lamprey gene sequences, argues for only one whole-genome duplication, with the remaining gene

copies produced afterward by smaller chromosomal duplications and translocations. Ocampo Daza's team rebuts this in the same paper, arguing that their two-duplication, fourfold-symmetry model is the more parsimonious account — parsimonious meaning it explains the same pattern using fewer separate assumed events, which is usually, though not always, the way to bet. The disagreement has not been settled by new data. It remains a live dispute about mechanism, layered under an already wide age estimate.

Whatever the precise mechanism, the family diversified, and every living jawed vertebrate carries some descendant of it. Ray-finned fish kept a version called isotocin. Birds, reptiles, and the mammals that never evolved a placenta carry mesotocin. Amphibians and most reptiles carry arginine vasotocin, differing from mammalian vasopressin by exactly one amino acid — a single spelling change in a nine-letter word, conserved since before mammals existed. Each name marks the same ancestral gene, inherited by a different branch and put to slightly different work.

Vasotocin is the best-mapped of the non-mammalian versions. A review led by Mario Penna at the Universidad de Chile, synthesizing decades of work by researchers including Boyd, Moore, Hollis, González, Smeets, and Vandesande and Dierickx, traces vasotocin-producing cells running from forebrain to hindbrain in newts, salamanders, frogs, and lizards — a wiring pattern conserved across the whole group. Newts named *Pleurodeles waltl* and *Taricha granulosa* carry it. So does the salamander *Plethodon shermani*. So do frogs across several genera: the European common frog *Rana temporaria*, the edible frog and marsh frog of the *Pelophylax* group, the American bullfrog, the common toad *Bufo bufo*, the Japanese toad *Bufo japonicus*, the African clawed frog *Xenopus laevis*, the cricket frog *Acris crepitans*, and the túngara frog from the opening scene. Two lizards, *Lacerta muralis* and *Acanthodactylus pardalis*, carry it as well. Across this list, vasotocin drives calling, clasping, and other pieces of courtship — but not as a fixed switch. In newts, the hormone only induces courtship behavior once the animal has already been exposed to the right level of a separate hormone, androgen; block that exposure and

vasotocin does much less. The molecule sets a possibility, not an inevitability.

Ray-finned fish took the ancestral gene in an oddly contested direction. Babak Ataei Mehr, Shawn Garner, and Bryan Neff at the University of Western Ontario injected isotocin — the fish paralogue of oxytocin, paralogue meaning a gene copy descended from the same ancestral duplication rather than borrowed from a different species — into ninety-two adult guppies, forty-six of each sex. Guppies given isotocin spent twenty-nine percent more time near other guppies than guppies given a chemical blocker of the same receptor, with no difference between the sexes. The same molecule turns up in zebrafish, in goldfish, and in the cooperatively breeding cichlid *Neolamprologus pulcher*. What makes the 2020 guppy result worth flagging honestly is that it cuts against most of what came before it: the authors note plainly that earlier fish studies had generally found no social effect from isotocin, or even an antisocial one. This is one of the first clean demonstrations of a pro-social effect, not a settled rule for fish in general.

Birds run the same gene as mesotocin, and in gregarious species it does something distinctive: it holds flocks together, not just pairs. James Goodson and colleagues at Indiana University found that blocking central mesotocin receptors in estrildid finches cut the time these birds spent with large groups and with familiar social partners. In the zebra finch, mesotocin acting on midbrain receptors is necessary for forming a pair bond in the first place; in the pinyon jay, it is implicated in broader social bonding; the common waxbill carries it too. The same molecule that binds a mammalian pair together here holds together an entire flock — an old tool reused at a different social scale, not a contradiction so much as a demonstration of how readily one signaling pathway gets redeployed.

Beyond the Backbone

The nonapeptide family is older than backbones. In *Caenorhabditis elegans*, a roundworm about a millimeter long with no brain to speak of, a peptide called nematocin acts through a receptor named NTR-1 to let the

worm learn an association between salt concentration and past experience — a capacity called gustatory associative learning, meaning the worm can change how it responds to a taste based on what happened the last time it met that taste. Isabel Beets and colleagues at KU Leuven, with a near-simultaneous and independent report from a second team, showed that worms missing nematocin or its receptor lose this learned response. Nematodes and vertebrates last shared a common ancestor before the protostome-deuterostome split, conventionally placed above 600 million years ago — pushing the functional origin of this signaling system deeper than any vertebrate fossil, into an ancestor whose nervous system had barely begun to specialize.

In the earthworm *Eisenia foetida*, the same family goes by annetocin, and it does something entirely different: injected in doses above five nanomoles, it triggers egg-laying — rotating body movements, a change in body shape, and the mucous secretion from the worm's clitellum that builds the cocoon holding its eggs. Chizu Oumi and colleagues at the Suntory Institute for Bioorganic Research described this in 1996. Earthworms and vertebrates, like nematodes and vertebrates, diverged before that same protostome-deuterostome split — a second, independent invertebrate lineage pointing at the same deep age from a completely different branch of the tree.

Snails carry a version called conopressin, and mollusks diverged from vertebrates at the same boundary, over 600 million years ago — a third independent phylum converging on the same depth. The great pond snail *Lymnaea stagnalis* shows something the vertebrate lineage no longer does: one single receptor, cloned by van Kesteren and colleagues at the Vrije Universiteit Amsterdam and named LSCPR, handles both the oxytocin-type reproductive signaling and the vasopressin-type metabolic signaling that vertebrates eventually split across two separate receptor families. The same peptide, chemically identified in the marine cone snail *Conus geographus* as well, still does both ancestral jobs at once in the pond snail — a plausible glimpse of what the system looked like before it specialized, and itself indirect evidence that the ancestral system had not yet split into separate paralogues this far back. Later work on the

nudibranch *Berghia stephanieae*, published in 2024, extended this picture rather than challenging it.

Laid end to end, the family tree runs from a millimeter-long worm tasting salt, through an earthworm laying eggs, through a snail running two jobs on one receptor, up through fish that shoal, frogs that call, finches that flock, to the mammalian pair that gave the whole family its popular name. Six branches, one ancestral trunk, and the same short peptide doing recognizably related work in every one of them: some version of drawing together, of committing reproductive or metabolic machinery to another individual's presence.

None of this makes the peptide a single, simple lever. Its job changes with the wiring it lands in — a pair bond in one species, a flock in another, an egg-laying reflex in a third. What survives across six hundred million years and half a dozen phyla is not a fixed behavior. It is a chemical vocabulary word, reused wherever a nervous system needs to say: stay close to this one. Not every ancient messenger speaks about closeness. Some speak about rank instead — and one of them is about to raise its claws.

Overtured if: a confirmed nonapeptide-family gene — a genuine ortholog of nematocin or vasotocin — turns up in a cnidarian or sponge genome, pushing the family's origin back past the last common bilaterian ancestor this chapter treats as its deepest confirmed boundary.

Sources

- **2021** · Ocampo Daza, Bergqvist, Larhammar, "The Evolution of Oxytocin and Vasotocin Receptor Genes in Jawed Vertebrates: A Clear Case for Gene Duplications Through Ancestral Whole-Genome Duplications" — [Frontiers in Endocrinology](#)
- **2017** · Penna et al., "Arginine Vasotocin, the Social Neuropeptide of Amphibians and Reptiles" — [Frontiers in Endocrinology](#)
- **2020** · Ataei Mehr, Garner, Neff, "Effect of isotocin on shoaling behaviour of the Guppy (*Poecilia reticulata*)" — [Animal Cognition, SpringerLink](#)

- **2009** · Goodson, Schrock, Klatt, Kabelik, Kingsbury, "Mesotocin and Nonapeptide Receptors Promote Estrildid Flocking Behavior" — [PubMed](#)
- **2012** · Beets et al., "Vasopressin/Oxytocin-Related Signaling Regulates Gustatory Associative Learning in *C. elegans*" — [Science](#)
- **1996 / 1995** · Oumi et al., "Annetocin, an annelid oxytocin-related peptide, induces egg-laying behavior in the earthworm, *Eisenia foetida*" — [PubMed](#); van Kesteren et al., "A Novel G Protein-Coupled Receptor Mediating Both Vasopressin- and Oxytocin-Like Functions of Lys-Conopressin in *Lymnaea stagnalis*" — [PubMed](#)

11. The Lobster's Posture

In a tank in the Department of Neurobiology at Harvard Medical School, a lobster hangs in the water, held just still enough for a fine needle. Edward Kravitz, Margaret Livingstone, and Ronald Harris-Warrick are not studying the lobster's claws or its shell. They are studying its posture, and they are about to change it with a few molecules of injected chemistry. A measured dose of serotonin goes into the circulation, and the lobster's body rises. Its claws spread open. It stands high on its walking legs, a posture that in a real tank means: I am not backing down. Switch the injection to a different messenger, octopamine, and the same animal's body drops low, its claws and limbs point forward, and it holds a posture that in a real tank means retreat.

Published in *Science* on the fourth of April, 1980, the result was not a vague mood shift. It was a whole-body postural program, switched on wholesale by one molecule and its opposite switched on by another. The team confirmed this was more than a surface impression by recording directly from the ventral nerve cord — the chain of nerve tissue running along a lobster's underside, doing the job a spinal cord does in a human, though built on a different architectural plan. Motoneurons, the nerve cells that give muscles their final marching orders, fired in opposing, coordinated patterns depending on which of the two chemicals was circulating. The posture was not incidental to the chemistry. It was the chemistry, expressed in muscle.

Serotonin and octopamine both belong to a chemical family called amines — small molecules built around a nitrogen atom, cheap for a nervous system to manufacture and well suited to broadcasting a signal rather than whispering it to one neighbor. The lobster experiment mattered because it showed these two amines acting as opposites, not merely as two items on a longer list of things a nervous system happens to make. One says: hold this posture. The other says: hold the opposite one.

The finding held up beyond the lab's original animals. Later work on the squat lobster *Munida quadrispina* — a species in the Anomura, family Galatheidae, a different branch of the decapod tree from the true lobsters and crayfish of the original paper — found the same opposition between the same two molecules. Decapod is simply the name for the crustacean group that includes lobsters, crabs, shrimp, and their squat-lobster relatives, walking on ten legs more or less. That the postural switch reappears across that group, rather than sitting in one species as a curiosity, is what turns a single striking experiment into a conserved piece of nervous-system architecture.

Octopamine has long carried a tidy label: the invertebrate's version of noradrenaline, the arousal-and-alertness messenger that runs the vertebrate fight-or-flight response. Gáspár Jékely, working at the Living Systems Institute in Exeter, showed the real picture is less tidy and more interesting. Comparing genomes across the invertebrate world — echinoderms and their relatives, mollusks and earthworm-type animals, arthropods and roundworms, and a small marine worm called *Meara stichopi* that sits on its own separate branch of the animal tree — Jékely found working copies of adrenaline-type receptors distributed across all of them, not replaced by octopamine receptors but coexisting alongside them.

That distribution places noradrenaline-type signaling in the last common ancestor of essentially all bilaterally symmetric animals, a lineage that predates the protostome-deuterostome split, conventionally dated above 600 million years ago. Octopamine did not evolve to replace noradrenaline in invertebrates, the way an older, tidier textbook line once had it. Both systems, and a close chemical relative called tyramine, appear to have run side by side since deep in the animal family tree, with different lineages leaning more heavily on one or the other. Nobody has published a rebuttal of this reconstruction; it stands as a correction to the older, simpler story rather than a live dispute.

Borrowed for an Argument

None of that is where the lobster experiment became famous. In his 2018 book *12 Rules for Life*, the psychologist Jordan Peterson, of the University of Toronto, took Kravitz's finding somewhere its authors never claimed to send it. Lobsters and humans, he argued, diverged roughly 350 million years ago; lobster dominance hierarchies run on serotonin in a way he treated as directly analogous to human status hierarchies; and this shared biology, in his reading, shows that human hierarchy is a natural, evolved fact rather than a purely cultural or political arrangement. On Britain's Channel 4 News, he framed the same logic as a matter of inevitability — continuity between how animals and humans organize their social structures being, on his account, simply what evolution predicts, not something culture invented from nothing. No new experiment backs any of this. It is a popular synthesis, built on the Kravitz lineage of lobster research and extended to humans by analogy.

Leonor Gonçalves, a neuroscientist at University College London, took the analogy apart in January 2018. Her first objection is architectural: the brain structures serotonin acts on in vertebrates are far more varied than anything in a lobster's nervous system, and a lobster does not have a centralized brain in the vertebrate sense at all — it runs on ganglia, clusters of nerve cells strung along the body rather than gathered into one organ, and it has no amygdala, the almond-shaped vertebrate structure most tied to fear and emotional weighting. Her second objection is directional: the same molecule points different ways in different animals. Lower serotonin correlates with less aggression in lobsters, but in humans, weaker communication between the amygdala and the frontal lobe under low serotonin tends to impair emotional control and increase aggression instead — the opposite direction of effect from the same chemical shortage.

Gonçalves adds a pharmacological point and a comparative one. Prozac, a drug that blocks serotonin from being reabsorbed at nerve endings, does the same molecular job in a lobster's nerve terminals as it does in a human brain, and does not make a lobster anything a reasonable

observer would call happier: chemical similarity is not psychological equivalence. And for a comparison animal exactly as distant from humans as the lobster, she points to honeybees, whose natural hierarchy is matriarchal — males are driven out of the hive and die shortly after mating. If biology alone dictated one correct social structure, it is not obvious why insects and crustaceans, running much of the same old chemical toolkit, should end up prescribing such different ones. Gonçalves places the lobster/human split at the same rough 350 million years Peterson uses, but frames the one confirmed shared inheritance from that ancestor differently: not nervous-system wiring, but a digestive-system landmark, that ancestor being essentially the point at which animals first evolved a gut.

PZ Myers, an evolutionary developmental biologist at the University of Minnesota Morris who writes the blog Pharyngula, went after the number itself. Across two posts in January 2018 — the second built on a full argument from an anonymous colleague, whose work Myers published and endorsed without reservation — Myers argued that Peterson's 350-million-year figure is wrong by a wide margin. Arthropods and chordates were already separate phyla by the Cambrian, roughly 550 million years ago, and molecular-clock estimates place the actual split closer to 800 million years — a discrepancy of somewhere between two hundred and four hundred and fifty million years, depending which end of the range is used for comparison.

Myers's deeper objection is about what running on serotonin is even evidence of. Serotonin is used by essentially every animal with bilateral symmetry, and by some plants, fungi, and amoebae besides, so finding it in both a lobster and a human says very little on its own — it is the chemical equivalent of noting that both animals use water. Lobsters, he stresses, do not have brains in the sense vertebrates do. And hierarchy's shape is not fixed even among humans' nearest relatives: bonobos, genetically close to both chimpanzees and humans, run a promiscuous, female-led social system on brain chemistry barely distinguishable from a chimpanzee's strictly male-dominance-ranked one. Shared chemistry

sets a capacity for social structure. It does not write the structure's rules; in humans especially, culture does that part of the job.

It is worth noticing that Gonçalves and Myers do not even agree with each other about where Peterson goes wrong. Gonçalves accepts his 350-million-year figure and argues the analogy fails anyway, on architecture and on direction of effect. Myers rejects the figure itself as an elementary error. Two named specialists, converging on the same verdict that the argument fails, by two different and only partly compatible routes — which is a healthier sign for the underlying science than if every critic had recited an identical objection.

The critique traveled well past professional biology. Nieve Brydges, writing as a science columnist for *Varsity*, the University of Cambridge's student newspaper, published a 2022 piece granting a roughly 500-million-year vertebrate/invertebrate split, with lobsters diversifying some 150 million years after that base divide, and still finding the argument broken on its own terms. Evolution, in her framing, is not a goal-directed climb toward some improved final form, the way a game character advances by leveling up. A hierarchy that survived natural selection is not thereby a template worth imitating, and cooperative, egalitarian social structures survive natural selection just as well as strictly ranked ones do. Biology, she concluded, supplies no license for treating any one arrangement as the uniquely natural one. Brydges is flagged here as a non-specialist voice, not a research source, but her column shows how far downstream of the original biology this argument traveled, and how many independent readers kept finding the same seams in it.

What actually survives all this scrutiny is narrower than the popular version, and more interesting for being narrower. Serotonin and octopamine really do set opposite postures in decapod crustaceans, wired down to the motoneuron. The amine-signaling toolkit really does trace back to the last common bilaterian ancestor, older than the animal kingdom's split into separate body plans. What does not survive is the leap from that toolkit to a verdict about human society. A shared

molecule, redeployed by different nervous systems for different jobs, is a story about evolutionary reuse — an old word borrowed into a new language, not the same sentence repeated. It sets what a nervous system is capable of. It does not set what a society is required to be.

The lobster's flexed posture lasts only as long as the serotonin does, and drops the moment a bigger claw enters the tank. Rank, even in a lobster, is a state, not a sentence. Further along the same nerve cord runs a different ancient amine, tuned not to posture but to pursuit — the molecule that carries this journey forward, from standing one's ground to wanting something in the first place.

Overtuned if: a decapod crustacean is found in which octopamine produces the dominant flexed posture and serotonin the subordinate one — reversing rather than repeating the Kravitz pattern and breaking the cross-species conservation the squat-lobster replication currently supports.

Sources

- **1980** · Livingstone, Harris-Warrick, Kravitz, "Serotonin and Octopamine Produce Opposite Postures in Lobsters" — [Science](#)
- **2017** · Jékely, "Ancient Coexistence of Norepinephrine, Tyramine, and Octopamine Signaling in Bilaterians" — [BMC Biology](#)
- **2018** · The Conversation, "Psychologist Jordan Peterson says lobsters help to explain why human hierarchies exist – do they?" (quoting Jordan Peterson and Leonor Gonçalves, UCL) — [The Conversation](#)
- **2018** · Myers, "Lobsters are not people" — [Pharyngula/FreethoughtBlogs](#); companion post, "Jordan Peterson does not understand evolution, nervous systems, or logic" — [Pharyngula/FreethoughtBlogs](#)
- **2022** · Brydges, "Botched biology: how does Jordan Peterson use science to prop up his (mis)beliefs?" — [Varsity](#)

12. The Vertebrate Bargain

The nerve cord that ran serotonin and octopamine through a lobster's rise and retreat carries a third ancient amine as well, tuned not to posture but to pursuit: dopamine. Where the nonapeptides bind one animal to another, and the lobster's amines set its rank for only as long as the chemistry lasts, this messenger has a different assignment — not what to feel about another creature, and not where to stand, but whether to go after something that has not happened yet.

A lamprey in a cold river has no jaw, no bone in its skull, and barely a vertebra to its name — just a flexible rod called a notochord and a ring of cartilage. It is, together with the hagfish, one of the two oldest lineages of backboned animal alive today, a living stand-in for what the earliest vertebrates were like before anything grew a hinge in its mouth. Cut open its small brain and something unexpected turns up: a fully wired basal ganglia, the deep-brain structure that decides which movement gets to happen right now, running on dopamine in essentially the same layout a human's does.

The basal ganglia work like a traffic controller for action, holding back competing impulses until one is cleared to go. Dopamine is the chemical that tunes the controller's threshold. Losing the dopamine-producing cells in this exact circuit is what causes Parkinson's disease in humans — the tremor, the stiffness, a body that no longer obeys smoothly. In the lamprey, researchers have mapped the same wiring: dopamine-producing cells in the regions equivalent to the mammalian substantia nigra and ventral tegmental area, connected through the same "direct" and "indirect" pathways that let a mammal's brain both start and stop a movement, gated by the same two receptor types, D1 and D2, acting like two different switches on the receiving side. The comparison was drawn together in a 2021 review, and named outright five years earlier in a paper called, without much modesty, "The Basal Ganglia Over 500 Million Years," published in *Current Biology* in 2016.

Lampreys and hagfish split from the lineage leading to jawed vertebrates roughly 560 million years ago, in the late Ediacaran, before anything alive had a jaw or a limb. That makes them the most basal living vertebrates available for comparison. Because their dopamine-driven basal ganglia already show the full mammalian-style architecture, the system's origin can be placed at least that far back — to the root of Vertebrata itself. Some uncertainty remains about how much of the lamprey's other brain wiring mirrors the mammalian arrangement exactly, versus representing a simpler ancestral state elaborated on later. The core dopamine architecture is not in question.

The circuitry is old. The chemistry beneath it is older still. Dopamine is built in two steps from tyrosine, an amino acid the cell already had on hand for other jobs: an enzyme called tyrosine hydroxylase converts it to L-DOPA — the same molecule given as medication to Parkinson's patients today — and a second enzyme, DOPA decarboxylase, finishes the conversion into dopamine itself.

That exact two-step recipe turns up in an animal with no backbone at all. In 2021, a team from Stockholm University — Rehab El-Shehawy, Sandra Luecke-Johansson, Anton Ribbenstedt, and Elena Gorokhova — confirmed both enzymes expressed in the gut wall of the water flea *Daphnia magna*, tracked the second enzyme's activity across the animal's molt cycle, and detected biologically active L-DOPA, though not free dopamine, inside the gut. The paper appeared in *mSystems* on 9 November 2021.

Daphnia is a crustacean, a protostome — on the opposite branch of the animal family tree from every vertebrate, which are deuterostomes. The simplest reading of two such distant lineages sharing an identical two-enzyme pathway is inheritance from a shared ancestor, not two separate inventions of the same chemistry. That would push the pathway's origin back before the protostome-deuterostome split — and exactly when that split happened is itself an open argument, sitting somewhere in the Neoproterozoic. Molecular-clock studies place it at roughly 642 to 761 million years ago, with one eighteen-gene analysis landing near 670

million; fossil-calibrated statistical models pull it in tighter, to about 558 to 575 million years ago. The disagreement between these two methods — "rocks versus clocks" — is reviewed candidly in a 2016 paper on the limits of precision in molecular timescales.

The chemistry was not idle wherever it came from. In the roundworm *Caenorhabditis elegans*, a dopamine system built from exactly eight neurons shapes foraging decisions, as reviewed in 2025 in Royal Society Open Science. In the fruit fly *Drosophila*, a cluster of dopamine neurons called PAM is required for a fly to learn that a smell predicts sugar or water — an appetitive association, meaning one built around approaching something good rather than fleeing something harmful — and fires when the reward itself lands, a result reported in a 2025 preprint. In *C. elegans* specifically, dopamine seems to act more as a diffuse chemical mood-setter, bathing many neurons loosely, than as the point-to-point wiring vertebrates use, so the word "circuit" fits this invertebrate case more loosely than it does a mammal's brain.

Keeping Score

The clearest demonstration that dopamine is not simply a pleasure chemical came from recording individual dopamine neurons in the brains of behaving monkeys. Wolfram Schultz, working with theorists Peter Dayan and P. Read Montague, found these neurons firing not just when a reward arrives, but in proportion to how wrong the animal's expectation had been. A reward better than expected produces a burst above the neuron's normal rate. A reward that lands exactly as expected produces nothing unusual. A reward that fails to arrive, when one was expected, produces a dip below baseline, timed precisely to the moment it should have shown up. Published in *Science* in March 1997, the pattern matched, almost term for term, the error signal used in a family of computer learning methods called temporal-difference learning, where a system updates its guess about the future based on exactly how far off its last guess turned out to be.

The same signal has since turned up well beyond one macaque laboratory: prediction-error-shaped dopamine activity has been traced through teleost fish and rodents, and picked up in humans by fMRI scans of the ventral striatum and ventral tegmental area during simple reward-guessing tasks. A computation first isolated in a monkey's midbrain turns out to be a vertebrate-wide habit, not a primate specialty.

Nothing about that computation needed new molecular parts. It ran on wiring already in place for hundreds of millions of years — the same tyrosine-to-dopamine pathway found in a water flea's gut, installed in a circuit built at the root of Vertebrata, and handed a new job: not moving a limb, but keeping score.

The prediction-error story, though, is not the whole story. Kent Berridge and Terry Robinson, at the University of Michigan, found that injecting amphetamine directly into a brain hub called the nucleus accumbens made rats want sucrose more — approaching it, working harder for it — without making them like it more, judged by the same orofacial taste-reactivity expressions, a distinctive tongue-and-mouth pattern used to measure hedonic pleasure in infant mammals. Wanting and liking, they argued, are separable, run through different circuits, and dopamine's main contribution is generating the first: the pull of what they called incentive salience, not the pleasure of the reward itself.

The original theory dates to 1998, in *Brain Research Reviews*; it was sharpened and re-titled outright as "The debate over dopamine's role in reward" in a 2007 update. Extended into an incentive-sensitization theory of addiction, the same split treats craving that keeps intensifying even as a drug's pleasure fades as exactly what runaway "wanting," unmoored from "liking," would be expected to produce — the account Robinson and Berridge were still defending, thirty years on, in a 2025 retrospective. The three-way split — whether dopamine is mainly about learning, liking, or wanting — remains a live, unresolved question, not one any side has closed.

The receptors reading dopamine's signal multiplied along the way. A whole-genome duplication — an event in which an organism's entire

genetic library is accidentally copied twice over, doubling every gene at once — occurred near the root of the vertebrate lineage roughly 450 million years ago, in the Ordovician, splitting one ancestral pair of dopamine-receptor genes into the four D1-like and several D2-like receptors carried by fish, birds, and mammals today, mapped in a 2013 study in *Molecular Biology and Evolution*. Not every lineage kept every copy afterward — mammals later lost one of the four, called D1C, and a different one, D1E, disappeared independently in several other branches. That the surviving system still functions the same way across that whole 450-million-year gap is not only a genetic inference: zebrafish dosed with morphine develop a place preference blocked by the opioid-blocker naloxone, and the same fish show conditioned preference for nicotine and ethanol persisting after three weeks of abstinence, reported in a 2013 study in *PLOS ONE*. Zebrafish and humans do not look alike, do not share a common ancestor for 450 million years, and do not build their reward systems the same way anatomically. Whatever computes reward in each of them still runs on the same borrowed alphabet.

From a jawless fish with no cortex at all to a small striped fish choosing the tank chamber where a drug once sat, the wager is the same wager: predict, act, and let one ancient molecule keep the score. The molecule stayed this constant across every one of these animals; what changed, in two lineages more than any other, was the size and shape of the brain built around it.

Overtured if: the TH/DDC dopamine-synthesis genes found in *Daphnia* turn out to be non-homologous to the vertebrate versions — independently evolved rather than inherited from a shared ancestor — converting what reads as one deep inheritance into two separate inventions of the same two-step chemistry.

Sources

- **2021** · El-Shehawy, R. et al., "Microbiota-Dependent and -Independent Production of l-Dopa in the Gut of *Daphnia magna*" — [PMC8577283](https://pubmed.ncbi.nlm.nih.gov/3577283/)

- **2016** · "Uncertainty in the Timing of Origin of Animals and the Limits of Precision in Molecular Timescales" — [PMC4651906](#)
- **1997** · Schultz, W., Dayan, P., Montague, P.R., "A Neural Substrate of Prediction and Reward" — [Science 275:1593](#)
- **1998** · Berridge, K.C., Robinson, T.E., "What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience?" — [Brain Research Reviews 28:309](#)
- **2016** · "The Basal Ganglia Over 500 Million Years" — [Current Biology30680-7.pdf](#))
- **2013** · "Evolution of Dopamine Receptor Genes of the D1 Class in Vertebrates" — [PMC3603308](#)

13. Big Brains, Ancient Molecules

In the late Paleocene and early Eocene, some 55 to 60 million years ago, a lineage of hoofed land mammals waded into shallow water, then swam, then never came back out – the archaeocetes, ancestors of every whale and dolphin now living. Something happened to their brains on that journey that has an equal in only one other lineage on Earth: primates, including, tens of millions of years later, humans.

Comparing brain size across such different animals runs into a problem immediately. A bigger body needs a bigger brain for reasons that have nothing to do with intelligence – more skin to feel, more muscle to command, more distance for a nerve signal to travel. In 1973, Harry Jerison, at UCLA, proposed a fix: the encephalization quotient, or EQ, which divides an animal's actual brain mass by the brain mass statistically "expected" for a body of its size. An EQ of 1 means exactly average for that body size. Higher than 1 means more brain than body size alone would predict.

The formula itself – brain mass divided by 0.12 times body mass raised to the two-thirds power – matters less than what it is built to capture: relative brain size, standardized so that a mouse and a whale can sit on the same scale at all.

On that scale, humans score higher than any other animal measured, with an EQ of roughly 7.4 to 7.8. The sperm whale, holder of the largest brain that has ever existed in any animal – some 8 to 9 kilograms, nearly six times the mass of a human brain's 1.3 to 1.4 kilograms – scores an EQ of only about 0.58, below the mammalian average. Absolute size and relative size point in opposite directions. Gerhard Roth and Ursula Dicke, reviewing the numbers at the University of Bremen in 2005, drew their comparison table from Jerison's original 1973 dataset together with a 1987 anatomical survey by Haug, using EQ as the necessary tool for comparing such different animals while naming its weakness in the same breath: they argued that raw brain size, and especially the number of cortical

neurons and the speed signals travel along them, tracks cognitive ability more closely than EQ does. EQ itself is, by the account of researchers who study it closely, an argued and unsettled proxy for "intelligence" — not a settled one.

Restricted to cetaceans alone, the expansion was neither steady nor uniform. Stephen Montgomery and colleagues at University College London reconstructed brain and body mass across a dated cetacean and hippopotamid family tree in 2013, recovering explicit EQ values at points along it: an early branch near 0.371, a later one near 0.702, and an ancestral baleen-whale node near 1.144. The increase happened in episodic jumps within particular branches, not as one smooth climb, and toothed whales and baleen whales do not even follow matching trajectories — an unresolved discrepancy the authors note directly, alongside evidence — drawn from separate work by Barton and Capellini, and by Hager and colleagues, published between 2011 and 2012 — that brain mass and body mass may be shaped by partly separate genes rather than one linked package.

What resulted looks nothing like a primate brain. Lori Marino, then at Emory University, described the cetacean cortex as reaching complexity by multiplication rather than specialization: a thin cortical unit repeated over and over across a highly folded surface, rather than the thicker, more differentiated modules that build a primate cortex. Her 2004 paper drew on four years of fossil work across 66 skulls, cross-checked against data from 144 more. It is her own framework for how cetacean complexity arises mechanically, not a mechanism proven against rival accounts of primate-style complexity — but no rebuttal of it has surfaced since.

Two lineages, one that stayed on land and one that returned to the sea, reached large, complex brains by different anatomical routes, at different paces, under different rules of proportion. What they did not do differently is the chemistry running through the result. The dopamine system built at the root of Vertebrata — the same tyrosine-to-dopamine pathway, the same ancient receptor families — is the one both a primate

brain and a cetacean brain inherited before either began building anything new on top of it.

Neither lineage needed to invent a new chemical vocabulary to build something unprecedented. A folded, multiplied cetacean cortex and a thickened, modular primate cortex are different pieces of hardware, arrived at on different schedules. Wired into both is the same handful of ancient signaling molecules, doing the same basic jobs they already did in a fish or an early mammal, just at a larger scale and in a new arrangement. Complexity, on this evidence, looks like a matter of how existing parts get organized and how many of them there are — not of sourcing new parts.

A third lineage makes the same point with a stranger anatomy still. Birds evolved cortex-equivalent processing using a structure organized nothing like the mammalian neocortex: a "nuclear" pallium made of clustered neuron groups, instead of the mammalian arrangement in layers. For roughly a century, the relevant bird-brain structures were named and treated as a primitive, reptile-like "striatum," as though the avian forebrain were mostly reflexive tissue with nothing resembling a cortex in it at all — until an international group, the Avian Brain Nomenclature Consortium, led by Erich Jarvis and Anton Reiner, used molecular markers, connectivity tracing, and electrophysiology to show these structures are true pallium, the evolutionary counterpart of cortex, and renamed them accordingly in 2005. One region, the nidopallium caudolaterale, is now treated as functionally equivalent to a mammalian prefrontal cortex. Birds and mammals last shared a common ancestor conventionally dated to around 320 million years ago, in the Carboniferous, at the split between the lineages leading to reptiles and birds on one side and to mammals on the other — meaning bird and mammal brains built convergent higher-order processing from independent starting anatomy, over comparably vast stretches of geological time.

Even the messenger molecule itself can be pressed into a genuinely new structural role. In the octopus visual system, researchers have identified

a dopamine receptor with no equivalent in vertebrate biology: an ionotropic receptor, a fast physical channel that opens on direct contact, where every vertebrate dopamine receptor instead works through the slower, cascade-based metabotropic route. The finding, which named the new receptor DopC1, was posted as a preprint in April 2025 and has not yet completed peer review — a caveat worth carrying alongside it. Cephalopods and vertebrates split at the same protostome-deuterostome divide discussed for dopamine's chemical origin, meaning the molecule itself is inherited, but this particular way of reading it is a separate, convergent invention layered on afterward — fitting company for an animal already described in the cephalopod-neuroscience literature as the most behaviorally advanced invertebrate known.

Ancestry and Reinvention

Reviews synthesizing work across several phylogenomics groups between 2013 and 2021 draw a clean line through all of this. The basic neural parts list — neurons, synapses, neuropeptides, and the front-to-back patterning program that lays out a nervous system's plan before it is built — was already present in the common ancestor shared by cnidarians and every bilaterally symmetric animal. Cnidarians, such as jellyfish and sea anemones, have only a diffuse nerve net and no brain at all, and they sit entirely outside the group containing every animal with left-right symmetry, which places this shared toolkit above 600 million years old.

Every large, complex brain built from those parts since — in ribbon worms, in annelids, in cephalopods, in arthropods, in vertebrates — is now read as an independent construction project, not a shared inheritance of "brain" as such: an example of homoplasy, the term for a trait that evolved more than once rather than being passed down from one common possessor. That reading replaced an older view in which shared complexity across these groups was itself taken as evidence of one deep common origin. The newer position, built on fossil and developmental gene-mapping evidence gathered through 2021, keeps the

molecular toolkit's single origin while giving up the idea that complex brains themselves share that same single origin.

That is the honest shape of the answer here. Every dopamine molecule moving through a human prefrontal cortex, a dolphin's folded neocortex, and a bird's nidopallium caudolaterale descends from the same billion-year-old chemistry, inherited once, from one source. The organs doing the thinking around that chemistry do not share that single ancestry. They were built, independently, at least three or four separate times over, by lineages that never had anything to work with but old molecules and a very long stretch of time. The molecule is exactly as old in a bird's brain as it is in a human one; only the architecture wrapped around it is new. Trace any of these borrowed molecules back far enough, past every brain they were ever built into, and they converge on something older than any brain at all: the single ancestral population that first wrote the alphabet they are made of.

Overtured if: cetacean and primate brain enlargement were shown to trace to a shared genetic driver already present in their last common ancestor — a Mesozoic mammal with neither trait — rather than to independent mutations arising separately in each lineage, converting a twice-told story of convergence into one shared inheritance.

Sources

- **1973** · Jerison, H.J., Evolution of the Brain and Intelligence — [ScienceDirect topic overview](#)
- **2004** · Marino, L., "Cetacean Brain Evolution: Multiplication Generates Complexity" — [eScholarship](#)
- **2013** · Montgomery, S.H. et al., "The Evolutionary History of Cetacean Brain and Body Size" — [Evolution](#) 67:3339
- **2005** · Roth, G., Dicke, U., "Evolution of the Brain and Intelligence" — [Trends in Cognitive Sciences](#) 9:250
- **2005** · Jarvis, E.D., Reiner, A. et al., "Avian Brains and a New Understanding of Vertebrate Brain Evolution" — [Nature Reviews Neuroscience](#) 6:151
- **2025** · "Novel dopaminergic neurotransmission in the Octopus visual system" — [bioRxiv](#)

14. The Unbroken Thread

Roughly 4.2 billion years ago, in the Hadean eon — by one recent estimate, about 300 million years after the Moon itself finished forming — a population of cells was already alive on Earth, and already keeping the alphabet this book has followed. Not a single cell. A population: several related lineages, trading genes back and forth, one of which happened to be the branch that left descendants running today. Every rival branch in that population fell silent long ago; only its own toolkit survived. It is called, by convention, LUCA — the last universal common ancestor — and the word "last" describes a fact about extinction. It does not claim LUCA was first, alone, or simple.

In a 2024 reconstruction led by Edmund Moody and Philip Donoghue at the University of Bristol, working with an interdisciplinary team of geologists, paleontologists, and phylogeneticists, LUCA's genome is estimated to have encoded something like 2,600 proteins — comparable to a modern bacterium — built from 399 gene families assigned a probabilistic likelihood of ancestral presence, up from the 355 families a more conservative 2016 estimate allowed. Among those genes: an estimated 19 for CRISPR, the bacterial immune system that fights off viruses. LUCA, in other words, was already old enough, and embedded enough in an ecosystem of its own, to have enemies.

The 4.2-billion-year date — with a 95% credibility interval running from 4.09 to 4.33 billion years — comes from comparing paralogous genes, pairs of genes within one genome that arose from an ancient duplication of a single ancestral gene, used as a molecular clock and calibrated against the independent record of microbial fossils and isotope chemistry. It was published in *Nature Ecology & Evolution* in July 2024, and profiled that November in *Quanta Magazine*.

Named scientists have named specific doubts about it. Patrick Forterre, former head of microbiology at the Institut Pasteur, who helped organize the very conference that coined the term "LUCA," calls the 4.2-billion-

year date very difficult to reconcile with how long Earth needed to cool after the Moon-forming impact, and doubts that the pace of early evolution can be pinned down as precisely as the reconstruction claims. Luis Delage, at Cinvestav Irapuato, accepts the early date but not the complexity: pointing to differences in membrane chemistry and DNA-replication machinery between bacteria and archaea, he argues LUCA was probably simpler than a genome of 2,600 proteins implies. Neither objection erases the finding. Both keep it honest.

Three independent lines of evidence converge on one ancestral population rather than several separate origins of life. The genetic code – the rulebook mapping the alphabet's three-letter words to amino acids – is nearly universal. The ribosome, the molecular machine that reads that code and strings amino acids into proteins, is homologous across every domain of life, built in part from RNA that still carries traces of an older, simpler world. And every organism builds its proteins from left-handed amino acids and its genetic backbone from right-handed sugars, a handedness that did not have to come out the same way twice. Each of these is what a statistician would call a high-dimensional coincidence: one arrangement, out of many that would have worked just as well chemically, that all of life somehow shares.

That reasoning was formalized by Carl Woese, at the University of Illinois, whose ribosomal-RNA phylogenetics rebuilt the root of the tree of life, and carried forward by the reconstruction lineage running from Weiss and colleagues in 2016 to Moody's team in 2024. Woese's own 1998 model, published in PNAS, pictured a "genetic annealing" phase before LUCA: early lineages trading genes so promiscuously, and mutating so fast, that nothing resembling a stable species yet existed, until the exchange slowed and today's three domains – Bacteria, Archaea, and Eukaryota – annealed into separate, vertically inherited lines. A 2003 follow-up attempted to reconstruct that ancestor's actual gene content directly from comparative genomics. The fact of universal conservation is not contested by anyone working in this field. Woese's specific mechanism for how that universality was reached remains one candidate among

several, because population dynamics from before LUCA are far harder to test than LUCA's own genome content is.

What Would Break It

A four-billion-year story is only worth trusting if it says plainly what would break it. Paul Davies, at Arizona State University, built on an earlier proposal by Carol Cleland and Shelley Copley to describe the clearest test: a "shadow biosphere," microbial life using a genuinely different biochemistry — different handedness, no shared genetic code — that arose separately from the LUCA lineage and has simply gone undetected because standard microbiology is tuned to recognize only the biochemistry life already has. Cleland and Copley's 2005 paper first proposed hunting for such "shadow microbes"; Davies and coauthors generalized the idea into a full search program in 2009, proposing concrete strategies such as surveying extreme environments — deep-sea hydrothermal vents running at 300 to 400 degrees Fahrenheit among them — for chemistry matching nothing already known.

Nothing has been found. Every sample examined so far falls inside the one known biochemistry, the one alphabet this book has followed from its first chiral molecule onward. Davies himself frames what a genuine find would mean: "the biggest sensation in biology since Darwin." That a scientist would reach for a comparison that size, for a search that has so far come up empty, is itself an honest admission of how much the single-origin picture currently rests on the absence of a counterexample, not on proof that none exists.

The test runs in the other direction too — not only whether an independent origin has been found, but whether one could now be built. In December 2024, thirty-eight scientists from nine countries, including two Nobel laureates, published a joint call in *Science*, backed by a roughly 300-page independent technical report, for a moratorium and public dialogue on research aimed at creating "mirror bacteria": organisms built from D-amino-acid proteins and L-nucleic acids, every molecule's handedness reversed from every organism alive today. Their concern was

specific — such an organism could evade immune recognition in humans, animals, and plants simultaneously, because every immune system alive has only ever learned to recognize the one-handed chemistry shared by its own body and its pathogens. The technology itself is judged to be at least a decade away. No mirror organism exists in nature or in any laboratory, and the report notes that active efforts toward building one had already been halted.

That mirror life is treated as an engineering problem, not a chemical impossibility, is itself informative. It means the universal left-handedness inherited from LUCA was never the only workable answer chemistry could have produced. It was what biologists call a frozen accident: an early, contingent outcome that locked in because everything built afterward had to match it, not a law nature was required to arrive at. A frozen accident is still a real, causally binding fact about the world. It is simply not a necessary one.


The most concrete version of this test is written into the chemistry itself. A genuine excess of D-amino acids — not the small, non-biological racemic imbalances already observed in a few meteorites, but a real, biological-strength excess — found in any future sample would count as strong evidence of life unrelated to Earth's own L-amino-acid lineage. A 2020 review in *Chemical Reviews* found that most amino acids formed abiotically in meteorites are racemic, with no significant D-excess yet identified in any extraterrestrial sample; the small excesses that do appear in some primitive meteorites match terrestrial biology's own left-handed bias and are attributed to non-biological chemistry, not to D-dominant life elsewhere. D-amino acids do occur in nature, but only in narrow, non-structural roles — cross-links in bacterial cell walls, a handful of peptide antibiotics — never as the backbone of a whole organism's proteins. Any future chiral detection, the same literature cautions, would still need independent confirmation before being accepted as a biosignature, since not every chemical imbalance in nature has a biological cause.

What the evidence supports, honestly, is not one first molecule, spotlighted and singular, striking lucky in some warm pool or vent. It is one population: many related, competing chemistries, most of which left no descendants, one of which happened to be the branch that did, at a moment currently dated to about 4.2 billion years ago and disputed by name. Every organism since — the water flea, the lamprey, the zebrafish, the octopus, the sperm whale, the human being doing the measuring — is a later paragraph in a story whose first paragraph is a population, not a point.

The molecules do not know any of this. The dopamine now computing a reward-prediction error in a human brain does not remember the water flea whose gut wall first ran the same two-step chemistry, and it will not remember the next creature to carry it either. Every animal that has ever used this alphabet — to want something, to learn from being wrong about it, to bond, to posture, to think — has died, or will. The alphabet is still here, being read by another mouth, in another body, for the first time again. It does not need to remember its carriers, because it survives whether or not any of them do. That is the thread this book has followed across four billion years: not the immortality of the self, but of the sentence — the alphabet outliving every creature that ever spelled a feeling with it.

One tree, the same messengers

Where each signalling molecule sits on the tree of life. Branch lengths are schematic.



One tree of life, and where each messenger sits on it.

Overtured if: a genuine, biological-strength D-amino-acid excess is confirmed in any living organism, fossil, or extraterrestrial sample – proof of a shadow biosphere or a second, independent origin of life – converting "one ancestral population, so far as anyone has found" into two.

Sources

- **2024** · Moody, E.R.R. et al., "The nature of the last universal common ancestor and its impact on the early Earth system" – [Nature Ecology & Evolution](#)
- **2024** · Lambert, J., "All Life on Earth Today Descended From a Single Cell. Meet LUCA." – [Quanta Magazine](#)
- **1998** · Woese, C.R., "The Universal Ancestor" – [PNAS 95:6854](#)
- **2009** · Davies, P.C.W. et al., "Signatures of a Shadow Biosphere" – [PubMed 19292603](#)
- **2024** · "Confronting risks of mirror life" – [Science, DOI 10.1126/science.ads9158](#)
- **2020** · "The Search for Chiral Asymmetry as a Potential Biosignature in our Solar System" – [Chemical Reviews](#)