

Copies Before Originals

Convergence, Admissibility, and the Explanatory Status of the First Instance

Flyxion

July 2026

Abstract

Most accounts of similarity begin with descent: two things resemble each other because one was copied from the other, or both descend from a common source. This paper argues that shared ancestry, even where fully present and well documented, frequently does no explanatory work in accounting for convergent structure. The clearest demonstration is not a case of absent lineage but of screened-off lineage: twelve genetically identical populations of *Escherichia coli*, descended from a single ancestor, converge on similar adaptive outcomes for reasons unrelated to that shared ancestry, while one population's divergence is explained not by any difference in lineage but by its own accumulated history within an otherwise identical field of constraint. Extending this distinction — between historical origin, the first appearance of a structure in time, and explanatory origin, the constraint structure that makes its appearance and recurrence possible — across deeper evolutionary distance (independent carcinization in decapod crustaceans), intellectual history (the near-simultaneous, uncorrelated arrival of Darwin and Wallace at natural selection), and technological history (multiple independent origins of agriculture), the paper develops an Admissibility-Origin Theorem: convergence without shared lineage is a predictable consequence of a sufficiently narrow admissibility field, jointly with a trajectory's own position within it, not an anomaly requiring special explanation.

Mathematical convergence is deliberately excluded from this evidence and shown to belong to a different modal category entirely — necessity rather than likelihood — a distinction the paper states formally in order to avoid an equivocation its argument would otherwise invite. The paper closes by tracing the consequences of separating historical from explanatory origin for originality, independent invention, and authorship, using the calculus priority dispute between Newton and Leibniz to show how much institutional energy can be spent adjudicating historical origin in cases where it explains almost nothing, and supplies an explicit criterion for when the ordinary genealogical explanation remains the correct one. The original, on the resulting account, retains its historical status without retaining explanatory priority over the very recurrence it is so often credited with causing.

Contents

1. The Historical Origin Assumption	2
2. Two Routes to the Same Conclusion	4
2.4 Convergence and Contingency Revisited	6
3. A Case That Does Not Belong With the Others	8
4. Consequences	9
5. Objections and Replies	11
6. Conclusion	13
References	15

1. The Historical Origin Assumption

Most accounts of similarity begin with descent. Two things resemble each other, on the standard view, because one was copied from the other, or both were copied from a common source. A manuscript resembles its exemplar because a scribe transcribed it. A language resembles its parent because speakers inherited it. A species resembles its ancestor because reproduction transmitted its form. Resemblance, on this view, is evidence of a causal history of transmission, and the strength of the resemblance is read as evidence for the closeness of that history. This is not a fringe assumption. It is close to definitional in genealogy, textual criticism, comparative linguistics, and evolutionary biology alike: similarity is the trace left by a lineage, and to explain the similarity is to reconstruct the lineage that produced it.

This paper argues that lineage, while frequently present, is not what does the explanatory work the standard view assigns to it. Something else is doing that work, and lineage is often only its most visible symptom.

The clearest demonstration is not a case where lineage is absent, but a case where lineage is present, fully documented, and still explanatorily beside the point. In 1988, Richard Lenski founded twelve populations of *Escherichia coli* from a single ancestral strain and placed them in an identical glucose-limited medium that also contained citrate, a carbon source ordinary *E. coli* cannot metabolize under oxic conditions (Blount, Borland, and Lenski 2008). All twelve populations share a common ancestor in the strict, literal sense the standard view requires. For more than thirty thousand generations, none of them evolved the ability to use the citrate sitting in their own growth medium, despite each population testing billions of mutations across that span. Then, in one population and one only, a citrate-using variant appeared, traced to earlier mutations that had done nothing detectable on their own except quietly make the later step accessible.

Ask what explains the pattern across the other eleven populations — the broad, repeated convergence on similar adaptive trajectories that the experiment's long-running fitness data document — and the common ancestor from 1988 is the wrong place to look. By the time the convergent outcomes appear, each population has accumulated tens of thousands of generations of independent, unshared mutations; whatever the founding strain contributed, it is not what is doing the work of explaining why lineages that have not exchanged genetic material in decades keep arriving at recognizably similar solutions. The ancestor

is temporally real and genealogically undeniable. It is also, for this explanatory purpose, screened off: present in the history, absent from the reason. What is doing the work is the shared constraint each population is independently subjected to — the same medium, the same limiting resource, the same selective pressure — applied twelve times over to twelve populations that, after the first few generations, might as well be strangers to one another.

The one population that diverges is, if anything, the more instructive case. Its citrate-using variant did not arise because that lineage was somehow more closely related to some hidden source of the trait. It arose because a sequence of earlier, ordinary mutations — themselves unremarkable, most likely neutral or nearly so at the time — happened to leave that population, and only that population, standing somewhere the trait was reachable from. (The precise interpretation of this contingency remains actively debated: Van Hofwegen, Hasty, and colleagues, and separately Maisnier-Patin and Roth, have questioned whether Lenski's team's original framing understates how readily the citrate-using phenotype can arise under related conditions, and Lenski's team has defended its original interpretation in response. The dispute concerns how rare and how history-dependent the outcome really is, not whether the divergence occurred.) What the disagreement does not touch is the more basic structure of the case: eleven populations converge under a shared constraint, without the convergence being explained by their shared ancestry; one population reaches a different outcome, without that divergence being explained by any difference in ancestry either, since all twelve started identically. Ancestry is constant across all twelve populations. The outcome is not. Whatever is doing the explanatory work in both the convergence and the exception, it is not ancestry.

Call the first appearance of a structure, in time, its **historical origin**, O_H . Call whatever structure of constraints makes that appearance possible — and, typically, makes its recurrence elsewhere possible too — its **explanatory origin**, O_E . The Lenski populations make the distinction unusually sharp because O_H is, in the strictest sense available, identical across all twelve lineages: one ancestral strain, one founding event, one date. And yet O_H explains almost nothing about the pattern of outcomes thirty thousand generations later. O_E — the shared medium, the shared limiting resource, the shared selective structure independently imposed on each lineage — explains a great deal, including, in its more fine-grained form, why one lineage found a path the other eleven did not.

This is not a claim that historical origins do not exist, or that questions of priority and firstness are meaningless. Someone was first to arrive at citrate metabolism in the Ara-3 population, just as someone is first to observe any convergent trait, first to publish any independently-discoverable proof, first to build any independently-inventable tool. The claim is narrower and, for that reason, harder to dismiss: that being first is a fact about timing, not a fact that explains the structure itself. O_H marks *when* something appeared. O_E explains *why it was available to appear at all*, there or anywhere else the same constraints happen to apply. Confusing the two — treating the first arrival as though its priority were doing explanatory work that in fact belongs to the field of constraints it arrived within — is the mistake this paper is aimed at.

Admissibility-Origin Theorem. Let \mathcal{A} be an admissibility field over a space of possible structures X , and let $\tau_1, \tau_2, \dots, \tau_n$ be trajectories independently constrained by \mathcal{A} , with no admissible path from any τ_i to any τ_j for $i \neq j$ over the relevant span. If \mathcal{A} favors a narrow region of X , the trajectories may converge on structurally similar states despite this absence of mutual reachability — and, as the exception in a population subjected to the identical field demonstrates, may fail to converge when a trajectory's own history alters which region of \mathcal{A} is locally reachable to it. In neither case does the historical origin of the convergent structure explain its recurrence. The explanatory origin lies in the admissibility structure that makes the outcome reachable, together with the local position of a trajectory

within that structure: the historical origin is only ever the first point at which that joint structure happened to become visible.

Stated geometrically, what the theorem describes is two notions of distance coming apart. Let d_H measure historical distance — how far apart two trajectories are in shared ancestry, transmission, or contact — and let d_A measure admissibility distance — how far apart the states they arrive at are within the space \mathcal{A} constrains. The cases this paper is built on are exactly the cases where

$$d_H(\tau_i, \tau_j) \rightarrow \infty \quad \text{while} \quad d_A(\tau_i, \tau_j) \rightarrow 0.$$

Historical distance and admissibility distance are independent quantities, not two readings of the same fact. Ordinary intuition assumes they move together — that things far apart in history are far apart in kind, and things close in kind must be close in history — and the argument of this paper is, in one sentence, that this assumption is not guaranteed and frequently false.

The remainder of this paper is an attempt to take that theorem seriously rather than merely stating it. Section 2 works through the strongest evidence for it, beginning with the finer structure of the Lenski case and extending to convergence at greater evolutionary and cultural distance. Section 3 addresses the one class of case that looks similar but is not — mathematical convergence — and explains why it cannot be folded into the same argument without weakening both. The sections that follow ask what happens to concepts like originality, invention, and authorship once O_H and O_E are no longer treated as the same thing.

2. Two Routes to the Same Conclusion

The Lenski populations demonstrate the theorem by a route that might be called *screening-off*: a shared ancestor is present, but thirty thousand generations of independent mutation intervene between that ancestor and the convergent outcome, so that whatever the ancestor contributed is no longer doing the explanatory work. This is not the only route available. A second, more familiar one is available wherever the candidate lineages are separated by enough evolutionary or historical distance that a transmitted origin is not merely explanatorily inert but essentially absent — cases where O_H , if it can be located at all, sits so far back and is so unlike the convergent structures it supposedly explains that treating it as their explanation would be implausible on its face. Call this route *deep absence*, in contrast to Lenski’s *screening-off*. Both routes converge on the same conclusion, by different arguments, and having both available makes the case considerably stronger than either alone.

Carcinization. The clearest biological instance of deep absence is not the vertebrate and cephalopod eye, frequently cited but genetically and developmentally more entangled than the popular version of the story suggests. It is carcinization: the repeated evolution of a crab-like body plan — a broad, flattened, dorsoventrally compressed carapace, a reduced abdomen folded beneath the body, a widened sternum — across decapod crustaceans. The term was coined in 1916 by the zoologist L. A. Borradaile, who described the pattern as nature’s many attempts to arrive at a crab. Current phylogenetic work identifies at least five independent originations of the crab-like form within Decapoda, spanning both true crabs and several lineages of “false crabs,” such as porcelain crabs and king crabs, that are not closely related to true crabs at all despite converging on much the same body (Wolfe et al. 2021). The relevant point for the argument here is not the joke value of the phenomenon — carcinization has become something of an internet meme in its own

right — but its evidential structure: five or more lineages, separated deep enough in the decapod phylogeny that their most recent shared ancestor was manifestly not crab-shaped, independently arrive at recognizably the same solution.

This case deserves the same honesty the Lenski case was given. Not every one of the five-plus originations is equally clean. Some researchers have raised the possibility that a subset of carcinization events may reflect shared ancestral developmental machinery — a common underlying toolkit inherited from a more distant ancestor and merely activated similarly in separate lineages — rather than fully independent invention of the relevant developmental pathway from nothing (Wikipedia’s summary of the current state of debate reflects this; the primary phylogenetic literature, e.g. Wolfe et al. 2021, treats the question of how independent the underlying mechanisms are as still partly open). This qualification does not damage the argument; if anything it sharpens it in a useful direction. Even granting some shared deep machinery, the researchers who have studied carcinization most closely explicitly reject reading it as convergence toward a single “ideal” crab form that evolution is somehow aiming at — the various carcinized lineages differ in which specific features they emphasize and by which developmental routes they arrive at a broadly similar outcome. The pattern is convergence onto a *region* of a constrained space, arrived at by more than one path, not convergence onto a single privileged point. That is precisely the admissibility-field picture, and precisely not a teleological one: \mathcal{A} favors a region because that region satisfies shared functional constraints — protection, maneuverability in a benthic environment — not because the field is aimed at crabs as a destination.

The field’s structure also becomes visible in reverse. Crab-like body plans have been lost — reverted to a more elongate, non-crab form — independently at least seven times, a process researchers call decarcinization. This is the carcinization literature’s own version of the Ara-3 exception: evidence that the admissibility field is not a one-way ratchet toward a fixed outcome but a structured space that trajectories can also exit, when a lineage’s own history opens a path back out of the region most others remain confined to. A theory of convergence that could only explain arrival and never departure would be exactly the kind of naive doctrine this paper is trying to avoid; carcinization and decarcinization together supply the same convergence-and-exception structure Lenski’s experiment supplied, this time across genuinely deep phylogenetic distance rather than within a single controlled lineage.

Darwin and Wallace. The clearest instance of deep absence outside biology is not a case of technology or folklore, both of which are vulnerable to the objection that some faint chain of diffusion might always, in principle, be found. It is the near-simultaneous, fully independent arrival of Charles Darwin and Alfred Russel Wallace at the same theory of evolution by natural selection. Darwin had been developing the theory privately since the late 1830s, sharing it only with a small circle of correspondents, when in June 1858 he received a manuscript from Wallace — written from Ternate in the Malay Archipelago, where Wallace was collecting specimens with no access to Darwin’s unpublished work — outlining a theory of natural selection close enough to Darwin’s own that Darwin’s friends Charles Lyell and Joseph Hooker arranged for extracts of Darwin’s earlier manuscripts and Wallace’s essay to be read jointly before the Linnean Society of London on the 1st of July, 1858. Neither man was present at the reading. There is no plausible transmission event to appeal to: Wallace had not seen Darwin’s unpublished theory, and Darwin had not seen Wallace’s reasoning before Wallace’s letter arrived. What both men shared was not a lineage but a field — the same body of accumulating evidence in natural history, the same Malthusian argument about population pressure that both later credited as a catalyst, the same broader scientific environment constraining which theories of biological change were, at that moment, reachable at all. Two independent trajectories through that shared

constraint space arrived, within weeks of each other, at structurally the same theory.

Independent agriculture. The broadest instance of deep absence, and the hardest to attribute to any residual diffusion, is the domestication of plants and animals. The archaeological and genetic consensus identifies multiple, geographically separated centers of agricultural origin — including the Fertile Crescent, China, Mesoamerica, the Andes and Amazonia, and New Guinea — each domesticating locally available species independently, without contact with the others, across a span of several thousand years beginning in the early Holocene and extending, for some centers, well into the mid-Holocene (Diamond 1997; Smith 1998). No single act of invention, transmitted outward, explains the pattern. What explains it is a shared constraint that arrived, at different times but by a common cause, across unconnected populations: post-glacial climate stabilization, changing resource availability, and comparable pressures on hunter-gatherer subsistence strategies, constraining what was reachable from many different starting points as each, in its own time, encountered similar conditions. Agriculture is, in this sense, carcinization and the Darwin–Wallace case scaled up to the level of entire technological and economic systems: not one lineage of invention branching outward, but several independent trajectories entering the same admissible region because the region, not any one trajectory’s history, is what made the destination reachable.

2.4 Convergence and Contingency Revisited

The Lenski experiment did not arise in a vacuum. It is worth being explicit about the debate it entered, since a reader familiar with that debate might otherwise spend the rest of this paper mapping its argument onto positions it does not hold. In *Wonderful Life* (1989), the paleontologist Stephen Jay Gould argued from the Cambrian Burgess Shale fauna that evolutionary history is radically contingent. Replay the tape from any sufficiently early point, he argued, and small, inconsequential differences would cascade into an entirely different biological world — one with nothing resembling the outcomes actually observed. In *Life’s Solution* (2003), Simon Conway Morris argued the opposite. Marshaling an extensive catalog of convergent traits, including cases much like carcinization, he concluded that evolutionary outcomes are considerably more constrained and predictable than Gould allowed; some readings of his argument extend this toward the claim that something like intelligence, or even humanity, was close to inevitable. The two positions are usually presented as a binary: either history is contingent, and convergence should be rare and unremarkable when it occurs, or evolution is constrained toward predictable outcomes, and Conway Morris’s reading of convergence as pervasive and directional should hold.

This paper’s argument is not a restatement of either side, and it is worth saying plainly why. Gould’s question was whether replaying history would reproduce the same outcomes. Conway Morris’s question was whether certain outcomes are, in some strong sense, inevitable once evolution is underway at all. The question this paper has actually been asking is narrower and more tractable than either: not whether convergence is the rule or the exception in general, but under what conditions a given admissibility field \mathcal{A} produces convergent outcomes across independent trajectories, and under what conditions the same field permits divergence instead. This is closer to the position the field has itself moved toward since Gould and Conway Morris staged their disagreement. Losos (2017) offers the fullest book-length treatment of this reframing, surveying the same experimental turn in evolutionary biology — the Lenski population among its central examples — that lets the question be asked empirically rather than only rhetorically. Blount, Lenski, and Losos (2018), reviewing several decades of subsequent experimental and observational work on

the question, conclude neither that evolution is predominantly contingent nor that it is predominantly convergent, but that both patterns occur regularly, with the more productive question being what determines which one a given case will show.

That reframing can be stated formally, and doing so makes explicit what this paper has been assuming throughout. Rather than asking whether convergence occurs — a yes-or-no question, and for that reason not a very informative one — the relevant quantity is

$$P(\text{convergence}) = F(\mathcal{A}, x_0),$$

a function of the admissibility field's structure together with a trajectory's own position, x_0 , within it. The Lenski populations make this precise in a way few other cases can. All twelve share the same \mathcal{A} : an identical founding strain, an identical medium, an identical selective structure. What differs, by the time the relevant window opens more than thirty thousand generations in, is x_0 — the accumulated, unshared mutational history each population has independently built up by that point. Eleven populations reach the window with a value of x_0 from which the citrate-using region of \mathcal{A} is not accessible. One reaches it with a value of x_0 , shaped by earlier potentiating mutations, from which that region is accessible. Same field. Different reachable region. No contradiction, and no need to decide whether evolution “is” contingent or “is” convergent in general, since the honest answer is that the field determines a *distribution* over outcomes conditional on position, not a single verdict independent of it. Convergence, on this view, is a property of admissibility geometry — the shape and reachability of \mathcal{A} from a given position — rather than a property of lineage geometry, which is exactly why the Gould–Conway Morris debate, framed entirely in the vocabulary of lineage and history, could not settle it.

It is worth being equally explicit about what this paper is not claiming, given how easily convergence arguments slide into stronger ones. Conway Morris's own writing moves, in places, toward the suggestion that pervasive convergence points to a deeper purposiveness in the evolutionary process — a reading he has connected explicitly to theological argument elsewhere in his work. Nothing in the admissibility-field account developed here supports, or requires, any such conclusion. \mathcal{A} is a formal structure of constraints, favoring some regions of outcome-space over others for describable, non-purposive reasons — shared physics, shared chemistry, shared functional demands on a benthic body plan, a shared selective medium in a flask. Explaining why independent trajectories are drawn toward the same region of such a structure is a claim about the shape of a constraint space, not a claim about design, direction, or destination. The theorem of Section 1 says that historical origin does not explain recurrence and that something else does. It does not say, and this paper does not intend, that the something else is trying to arrive anywhere in particular.

Across all three cases — a biological body plan reached from more than five separate phylogenetic starting points, a scientific theory reached by two scientists with no contact between their reasoning, and a technological transformation reached by populations with no contact with one another at all — the pattern is the same one the Lenski experiment demonstrated by a different route. O_H , the historical origin, is answerable in every case: Borradaile named carcinization in 1916, though the underlying events long predate him; Darwin and Wallace are individually identifiable; some particular population somewhere was first to cultivate some particular plant. None of these facts explain the pattern of recurrence. What explains it, in each case, is \mathcal{A} : a structure of constraints reachable, independently, from more than one starting point, favoring a region of outcome-space narrow enough that unconnected trajectories keep landing inside it.

3. A Case That Does Not Belong With the Others

There is a form of convergence without lineage more extreme than any considered so far, and it is tempting to reach for it as the strongest evidence available. Nobody copies the fact that two and two make four. Every competent reasoner who works through the relevant proof arrives at the Pythagorean relationship between the sides of a right triangle, and something close to this result was independently established in Babylonian, Chinese, Indian, and Greek mathematical traditions, separated by centuries and by any plausible route of transmission. If carcinization and the Darwin–Wallace case count as evidence that similarity does not require lineage, mathematical convergence looks like the cleanest case of all: total independence, total agreement, no contact required.

This is precisely the case this paper declines to use, and the reason is not that the phenomenon is unreal. It is that folding it into the same evidential category as Sections 1 and 2 would blur a distinction the argument depends on keeping sharp.

In every case considered so far, the admissibility field \mathcal{A} favors an outcome without requiring it. Carcinization is reached by more than five lineages, but crab-like forms have also been abandoned at least seven times; decarcinization is not a violation of the field, it is evidence that the field leaves room for trajectories to exit the favored region. The Ara-3 population found citrate metabolism precisely because its own mutational history opened access to a part of the field the other eleven populations' histories did not. In each case, the alternative to convergence — divergence, exception, a trajectory that lands somewhere else — is not merely unobserved. It is a live possibility the field itself permits, realized in exactly the cases examined for that reason. This is what makes these cases cases of *convergence* rather than *necessity*: the field assigns high likelihood to a region without assigning zero likelihood to everything outside it.

Mathematical truth admits no such counterpart. There is no population of reasoners, however constituted, for whom two and two make five under standard arithmetic, and no sense in which some rare enabling condition could open a path to that outcome the way a potentiating mutation opened a path to citrate metabolism. The relevant “alternative” is not merely unfavored by the field; it is not available at all, on pain of contradiction. Where biological and cultural convergence are best modeled as an admissibility field that concentrates probability or reachability on a favored region while leaving nonzero measure elsewhere — a highly biased attractor, not an exclusive one — mathematical necessity is the degenerate limit of that same picture, in which every alternative has been driven to zero. Convergence is a claim about where independent trajectories are likely to arrive. Necessity is a claim about where every coherent trajectory must arrive, because nothing else is a trajectory through the relevant space at all.

Convergence: $\mathcal{A}(x^*) \gg \mathcal{A}(x)$, $x \neq x^*$, $\mathcal{A}(x) > 0$ for some $x \neq x^*$.

Necessity: $\mathcal{A}(x^*) > 0$, $\mathcal{A}(x) = 0$ for all $x \neq x^*$.

This is not a difference of degree that a sufficiently strong constraint could bridge. It is a difference in what kind of claim is being made — nomological or statistical in the first case, logical in the second — and an essay that traded on the rhetorical strength of the mathematical case while actually building its argument from the biological and cultural cases would be quietly equivocating between them. A reader who noticed the equivocation would be right to distrust the rest of the essay's evidence as a result.

What mathematics does still share with the cases in Section 2, and this is worth stating precisely rather than discarding along with the rest of the analogy, is the same failure of O_H to explain what needs explaining. The Babylonian scribes who first recorded the

relationship encoded in the Pythagorean theorem did not make it true, any more than Darwin's priority over Wallace by a matter of weeks explains why natural selection was the theory reachable by both of them. What explains the theorem's truth, and its independent rediscovery in every tradition that developed the relevant geometric and arithmetic apparatus, is the structure of Euclidean space and the axioms from which the result follows — an explanatory origin, in the sense this paper has been using the term throughout, that happens to take the limiting, deterministic form of logical entailment rather than the probabilistic form of a biased attractor. $O_H \neq O_E$ holds here too. It simply holds for a different reason, and the paper is more honest, not less interesting, for saying so plainly rather than letting the two kinds of case blur into one undifferentiated pile of "convergence."

4. Consequences

The preceding sections established that $O_H \neq O_E$: that a historical origin, however well documented, does not by itself explain the recurrence of the structure it originated, and that an admissibility field frequently does. What remains is to ask what this costs, and what it buys, once it is taken seriously outside the laboratory and the fossil record — in the ordinary vocabulary of originality, invention, authorship, and credit, all of which are built, more or less explicitly, on the assumption this paper has been dismantling.

4.1 Originality Without Priority. The ordinary use of "original" conflates two things that the preceding sections have pulled apart: being first, and having genuinely done the thing. On the historical view, these collapse into one another almost by definition — to be original just is to be the source others copied from, and priority in time is what establishes the claim. Once O_H and O_E are separated, this collapse looks like an accident of a world in which independent arrival happens to be rare enough, in most everyday domains, that firstness and genuine engagement with a problem's real constraints tend to coincide. The Lenski populations, and carcinization, and Darwin and Wallace, are precisely the cases where they come apart, and what they suggest is that firstness was never the property doing the work. What was doing the work, in every case where the ordinary judgment of originality still feels earned, is that a trajectory actually traversed the relevant constraint structure — solved the problem, survived the selective pressure, derived the proof — rather than merely reproducing a description of someone else's traversal without doing the work the description records.

This gives originality a definition that does not depend on a temporal ordering between trajectories at all. A trajectory's arrival at x^* is original if it is reached by engaging the constraints of \mathcal{A} directly — by testing mutations against a selective environment, by reasoning from axioms to a proof, by confronting a technical problem with the tools genuinely available — rather than by copying an already-realized instance of x^* without that engagement. On this definition, Wallace's arrival at natural selection is fully original despite Darwin's small priority in developing the idea, because Wallace reached it by his own route through the same constraint structure, not by reading Darwin's unpublished notebooks. A student who reproduces a known proof of the Pythagorean theorem without working through the geometric reasoning has not done anything original, regardless of whether they are the ten-thousandth person to write those exact steps or, by some accident of isolation, the first. Originality, on this account, is a relation between a trajectory and a field. Priority is a relation between a trajectory and a clock. The two frequently travel together. They are not the same relation, and treating them as though they were is precisely the mistake this paper has been arguing against since Section 1.

4.2 Invention and Discovery. Nothing does more to obscure the distinction between

O_H and O_E than the assumption, built into ordinary talk of invention, that independent arrival at the same solution requires an explanation beyond the ordinary — coincidence, cryptomnesia, undisclosed contact, espionage. Patent systems have, in fact, had to build formal machinery to handle exactly this possibility, because it turns out not to be rare enough to leave to intuition: for most of the twentieth century, United States patent law operated on a first-to-invent standard precisely because two parties independently arriving at the same invention was common enough to require a dedicated legal procedure, the interference proceeding, for adjudicating whose claim would stand. That such a procedure had to exist, and had to be used with some regularity, is a quiet institutional admission of the thesis this paper is making explicitly: given a sufficiently constrained technical problem and enough independent parties working on it, more than one trajectory reaching the same solution is not an anomaly to be explained away. It is what a narrow, strongly-favored region of \mathcal{A} predicts.

Darwin and Wallace are the clearest instance already discussed, but the same pattern recurs wherever a technical or scientific field narrows the space of viable next steps enough that multiple researchers, working from a shared body of available results, arrive at a shared destination without arriving at it from each other. This should not be surprising once O_E is understood as the operative explanatory structure. The interesting question, for any given case of apparently independent invention, is not “how could this happen without one party copying the other,” which presupposes that convergence requires special explanation, but “what did the shared field look like, and how favored was the region both parties reached.” Where the field is narrow and the constraints are demanding, multiple discovery should be the expectation, not the puzzle.

4.3 Authorship and Attribution. If originality does not track priority, and independent invention is unsurprising rather than anomalous, a further question becomes unavoidable: what is authorship actually doing, as a category, if it is not identifying the explanatory source of the thing authored? The calculus priority dispute between Newton and Leibniz is instructive here precisely because it is not one of this paper’s evidentiary cases. Both mathematicians developed calculus independently in the latter decades of the seventeenth century, working from a shared body of antecedent results — Fermat’s methods for finding tangents and extrema, Barrow’s work relating tangent and area problems, Wallis’s work on infinite series — that was, to some degree, available to both. This partial sharing of antecedents is exactly why the case was set aside in Section 2: it is not clean evidence of convergence without any relevant transmission, since some of the relevant mathematical raw material had genuinely diffused to both parties. But the case is very well suited to the question this section is asking, because the dispute that followed consumed decades of bitter institutional energy — accusations of plagiarism, a Royal Society inquiry convened, awkwardly, under Newton’s own presidency, and a priority verdict that historians of mathematics now regard as having been decided more by institutional power than by the underlying facts (Hall 1980; Whiteside 1967–1981) — almost entirely over the question of O_H , who reached the result first and whether the other had seen it. Almost none of that energy was spent on the question that actually explains why calculus was discoverable by two mathematicians in the same few decades: the antecedent mathematical field itself, built up over the preceding half-century, had narrowed the space of viable next steps enough that the result was, in the sense this paper has been developing, close to reachable from multiple directions at once. The priority dispute is a monument to how much attention O_H can absorb precisely in the cases where it explains the least.

Authorship, on the account this suggests, is best understood not as a claim about explanatory origin at all, but as a social and institutional convention for allocating credit, responsibility, and reward — a genuinely useful convention, load-bearing for norms of

citation, compensation, and trust, and not one this paper is arguing against. What it is not is a reliable guide to why the authored thing exists or why it might recur elsewhere without the author's involvement. An author is very often the historical origin of a particular instance. Treating that historical fact as though it settled the explanatory question — as though the author's priority were what made the result true, or the invention reachable, or the trait viable — repeats, in the register of credit and reputation, the same error the rest of this paper has been diagnosing in the register of biology and cultural history.

4.4 When Lineage Still Wins. None of this licenses treating every instance of similarity as a candidate for admissibility-field explanation. Most similarity, in fact, is exactly what the historical view says it is: the trace of a transmission event, and the right explanation is genealogical, not structural. The question this paper's argument actually raises is not whether lineage explanations are usually wrong — they are usually right — but how to tell, in a given case, which kind of explanation applies. A rough criterion, in the spirit of a likelihood comparison between two competing explanations of an observed similarity, is to ask which hypothesis renders the resemblance more probable:

$$L(\text{Lineage}) \text{ versus } L(\mathcal{A}).$$

Where a transmission channel is demonstrable, or even merely plausible and unexcluded, $L(\text{Lineage})$ dominates and ordinary historical explanation remains the right tool. This is, on reflection, why internet memes were set aside in Section 2 rather than used as evidence: not because independent convergence in cultural production is impossible, but because in the overwhelming majority of cases that look like it, careful tracing — the kind of forensic reconstruction that dedicated meme-archival communities routinely perform, following timestamps, cross-posts, and screenshots back to a point of origin — recovers an actual transmission chain. Internet culture is, if anything, an unusually well-documented medium for lineage reconstruction, not a domain where lineage is hard to establish; the earlier decision to exclude memes from the evidence in Section 2 and the position taken here are the same position stated twice, once as a caution about weak evidence and once as a positive account of why the evidence is weak. Plagiarism detection, textual stemmatics — the reconstruction of a manuscript tradition's lineage from patterns of shared scribal error, a discipline built entirely on the premise that certain kinds of resemblance are only explicable through a transmission history — and the identification of forged or counterfeit goods all belong to this same category: domains where $L(\text{Lineage}) \gg L(\mathcal{A})$ as a matter of course, and where reaching for an admissibility-field explanation would not be a subtler account of the resemblance but a mistaken one.

The admissibility-field explanation earns its place only where this ordering reverses: where a transmission channel can be excluded, as with Darwin and Wallace, or rendered explanatorily irrelevant by the sheer span of independent history intervening, as with the Lenski populations, or where the same outcome recurs across a distance — phylogenetic, geographic, or temporal — wide enough that no plausible channel connects the instances at all, as with carcinization and independent agriculture. The criterion is not a license to prefer the more elaborate explanation. It is a discipline for using it only where the simpler one has already been ruled out on the evidence.

5. Objections and Replies

The argument of this paper has been built case by case, with objections handled inline as each case raised them: the indexed-distinction reduction in Section 1, the deep-homology qualification on carcinization in Section 2, the necessity/convergence distinction in Section

3. Four further objections are general enough that they apply to the paper's whole argument rather than to any one case, and are better answered directly than left for a reader to raise unassisted.

5.1 Isn't This Just Convergent Evolution? A reader with a background in biology may hear the argument of Sections 1 and 2 and conclude that the paper has done nothing more than restate convergent evolution in unfamiliar vocabulary. This mistakes the evidence for the thesis. Convergent evolution — carcinization, camera eyes, echolocation — is a *phenomenon*: independent lineages arriving at similar structures. The claim of this paper is not that this phenomenon occurs. Evolutionary biology has known that for well over a century, and needs no help from an essay on explanatory origin to keep believing it. The claim is about explanatory structure: that in cases like these, the first appearance of a structure does not explain its recurrence, and something else — an admissibility field, favoring a region of possibility space independent of any one trajectory's history — does. A biologist who accepts carcinization as thoroughly convergent and explains it in the ordinary way, through natural selection acting independently on separate lineages, is not disagreeing with this paper. Natural selection acting independently on separate lineages *is* an admissibility-field explanation, in every respect but name: it explains recurrence by appeal to a constraint structure reachable from more than one starting point, not by appeal to shared descent. What this paper adds is not a new biological claim. It is the observation that biology has been making this kind of explanatory move correctly for generations, while intellectual history, the history of technology, and the ordinary vocabulary of authorship and invention largely have not — and the argument of Sections 2 through 4 is an attempt to extend a form of explanation biology already trusts into domains that still default to lineage as though no alternative existed.

5.2 Doesn't \mathcal{A} Just Rename the Explanation? This is the strongest objection available, and it deserves a direct answer rather than a dismissive one. If “admissibility field” is simply a new label pasted over whatever historical origin used to explain, the entire apparatus of Sections 1 through 4 is decoration, not argument. The reply is that \mathcal{A} earns its place only if it generates predictions a lineage-only account does not, and it does. A lineage account predicts that similarity should track transmission: closely related or historically connected trajectories should resemble each other more than distant or unconnected ones, and a sufficiently thorough investigation should, in principle, turn up a channel of contact wherever real similarity is found. An admissibility account predicts something different: similarity should track position within the constraint structure, largely independent of genealogical or historical distance, and recurrence should appear wherever the same constraints recur, whether or not any transmission channel exists. These are not the same prediction wearing different words. The Lenski populations make the difference precise rather than rhetorical. All twelve populations are equally related to one another — the same founding strain, the same generation of separation — so a lineage account has no resource for explaining why eleven converge and one does not, since genealogical distance among the twelve is uniform by design. An admissibility account explains the split without difficulty: what varies is x_0 , position within the field, not degree of relatedness. Where the two accounts make the same prediction, as in most everyday cases of transmitted similarity, there is no need to invoke \mathcal{A} at all, and Section 4.4 says so explicitly. Where they diverge, as in Lenski, carcinization's independent origins, and Darwin and Wallace's simultaneous arrival, the admissibility account predicts what actually happened and the lineage account does not. A relabeling could not do that.

5.3 Is the Thesis Falsifiable? An admissibility claim, applied to a specific case, fails under three conditions, and the paper's own methodology in Sections 2 and 4 has been applying them throughout rather than stating them separately. It fails if a transmission

channel is subsequently discovered where none had been identified — had it emerged that Wallace saw Darwin’s private notebooks before writing from Ternate, the admissibility reading of that case would collapse into an ordinary, if embarrassing, case of undisclosed transmission. It fails if recurrence does not occur where the field predicts it should — if a structure is claimed to be strongly favored by a shared constraint, and many independent trajectories are exposed to that constraint for a comparable span without ever converging on it, the claim that the field favors that region is simply wrong, in the same way a claim about a biased attractor is wrong if the attractor turns out not to attract. And it fails, more mundanely than either, when purported convergence collapses into genealogy on closer investigation, which happens routinely in the actual practice of scholarship: cases initially reported as independent invention are not infrequently traced, on further archival work, to a forgotten correspondence, a shared teacher, or an overlooked prior contact, and when that happens the honest response is to withdraw the admissibility reading, not to defend it. The general thesis of this paper — that $O_H \neq O_E$ in at least some cases — is supported existentially, by the specific, individually checkable cases in Section 2, each of which could in principle have failed one of these three tests and did not. It is not offered as a universal law that no similarity is ever explained by lineage; Section 4.4 exists specifically to state when lineage remains the right explanation, and a thesis that predicted admissibility everywhere would be exactly the unfalsifiable overreach this paper has tried, throughout, not to become.

5.4 Why Distinguish O_H and O_E at All? A remaining skeptical move grants everything so far and asks why the distinction matters: historical origin and explanatory origin might simply be two aspects of one underlying process, and insisting on separating them could look like manufacturing a problem where ordinary language already gets by with a single word. The reply is that the distinction earns its keep exactly where disputes that looked unified turn out, once separated, to be about only one of the two things. The Newton–Leibniz controversy consumed decades of institutional energy establishing O_H — who reached the calculus first, who saw whose papers — while the question of O_E , why the calculus was reachable by two mathematicians working from a shared antecedent field in the same few decades, went almost entirely unaddressed by either side. Patent interference proceedings exist as standing legal machinery for exactly the same reason: two parties can both have genuinely invented the same thing, and the law has had to build a dedicated procedure for adjudicating O_H precisely because O_H and O_E do not automatically coincide often enough to be left to intuition. If the two collapsed into a single notion, as the ordinary use of “origin” suggests they should, none of these disputes would make sense as disputes: there would be nothing left to argue about once priority was settled, because priority and explanation would be the same fact under two names. That they are not the same fact — that a priority dispute can be settled definitively, as the Royal Society’s report claimed to settle the calculus dispute, while leaving the explanatory question completely untouched — is the clearest evidence available that the distinction was never a manufactured one. It was already doing work. This paper has only made it explicit.

6. Conclusion

The argument of this paper began with an assumption ordinary enough to be nearly invisible: that similarity is evidence of descent, and that to explain a resemblance is to trace the lineage that produced it. Twelve *E. coli* populations, sharing a single ancestor, made the assumption’s limits visible in miniature — not by lacking a lineage, but by demonstrating that the lineage they shared had stopped doing any explanatory work long before the

pattern it was supposed to explain appeared. Carcinization, arising independently across more than five branches of the decapod tree, and lost independently at least seven more times, showed the same structure at a distance too great for any lineage to bridge at all. Darwin and Wallace, reasoning in separate hemispheres from a shared body of accumulating evidence, showed it in the register of ideas rather than organisms. Independent agriculture showed it at the scale of entire subsistence systems. In every case, the question that mattered was not who or what came first, but what structure of constraint made the outcome reachable, more than once, from more than one starting point. That structure — the admissibility field, \mathcal{A} — is the paper's answer to the question its title asks. Mathematics showed where the same question changes character entirely, becoming a matter of necessity rather than likelihood, and the consequences traced in Section 4 showed what happens to originality, invention, and authorship once the historical origin of a thing is no longer mistaken for the reason it exists.

Section 5.4 already made the general shape of this operation explicit: the historical view treats first appearance and explanatory source as a single, undifferentiated category, and the admissibility view separates them. What is worth adding here, now that the whole argument is in view, is how much followed from that one separation. Four independent evidence classes, a modal exclusion, and four sections of consequences all trace back to the same refusal — the refusal to let two questions collapse into one merely because ordinary language has a single word, origin, for both of them.

It is worth returning, in closing, to the word the title actually uses. A copy, in the ordinary sense, is defined by its descent from an original: it exists because something else existed first and was reproduced. Most of this paper has not been about copies in that sense at all. It has been about convergence, recurrence, independent rediscovery — cases in which nothing was reproduced, because there was nothing available to reproduce from. And yet the outcomes in each case are copy-like in every respect that matters except the one the ordinary definition insists on: they resemble each other closely enough that, encountered without their histories attached, no observer would hesitate to call them copies of one another. The mistake the ordinary definition makes is to assume that this kind of resemblance always requires a source to have been copied. Sometimes what recurs is not a transmission. It is a possibility, realized more than once because the space of possibilities was never as open as the assumption of originality requires it to be. A citrate-metabolizing lineage, a crab-shaped body, a theory of natural selection, a cultivated grain: each looks, from the outside, like a copy of the others' kind. None of them copied anything. They arrived.

This is the sense in which copies come before originals, and it is worth stating without the provocation the title trades on, now that the argument behind it has been made in full. O_H , the historical origin, marks when a structure first appeared. O_E , the explanatory origin, is the field that made its appearance possible — there, and, under the same constraints, wherever else the same constraints happen to apply. The original deserves its historical status. It was first. What it does not automatically deserve, and what this paper has argued it is rarely entitled to, is explanatory priority over the very recurrence it is so often credited with causing. The recurrence of a structure is not explained by who or what entered a region of possibility first. It is explained by the field that made the region reachable to more than one trajectory at all — and the first arrival, however genuine its priority, is only ever the earliest place that field happened to become visible.

References

- [1] Blount, Zachary D., Christina Z. Borland, and Richard E. Lenski. 2008. "Historical Contingency and the Evolution of a Key Innovation in an Experimental Population of *Escherichia coli*." *Proceedings of the National Academy of Sciences* 105(23): 7899–7906.
- [2] Blount, Zachary D., Richard E. Lenski, and Jonathan B. Losos. 2018. "Contingency and Determinism in Evolution: Replaying Life's Tape." *Science* 362(6415): eaam5979.
- [3] Losos, Jonathan B. 2017. *Improbable Destinies: Fate, Chance, and the Future of Evolution*. New York: Riverhead Books.
- [4] Borradaile, L. A. 1916. "Crustacea. Part II. Porcellanopagurus: An Instance of Carcinization." *British Antarctic ("Terra Nova") Expedition, 1910–1913, Natural History Report, Zoology* 3: 111–126.
- [5] Wolfe, Joanna M., Javier Luque, and Heather D. Bracken-Grissom. 2021. "How to Become a Crab: Phenotypic Constraints on a Recurring Body Plan." *BioEssays* 43(5): 2100020.
- [6] Darwin, Charles. 1859. *On the Origin of Species*. London: John Murray.
- [7] Darwin, Charles, and Alfred Russel Wallace. 1858. "On the Tendency of Species to Form Varieties; and on the Perpetuation of Varieties and Species by Natural Means of Selection." *Journal of the Proceedings of the Linnean Society of London (Zoology)* 3(9): 45–62.
- [8] Gould, Stephen Jay. 1989. *Wonderful Life: The Burgess Shale and the Nature of History*. New York: W. W. Norton.
- [9] Conway Morris, Simon. 2003. *Life's Solution: Inevitable Humans in a Lonely Universe*. Cambridge: Cambridge University Press.
- [10] Diamond, Jared. 1997. *Guns, Germs, and Steel: The Fates of Human Societies*. New York: W. W. Norton.
- [11] Smith, Bruce D. 1998. *The Emergence of Agriculture*. New York: Scientific American Library.
- [12] Hall, A. Rupert. 1980. *Philosophers at War: The Quarrel Between Newton and Leibniz*. Cambridge: Cambridge University Press.
- [13] Whiteside, D. T., ed. 1967–1981. *The Mathematical Papers of Isaac Newton*. 8 volumes. Cambridge: Cambridge University Press.