

The Cave-Design Error

Reference Classes, Comparative Likelihoods, and the Burden of Agent Inference

Phil Stilwell

Independent Scholar

ORCID: [0009-0009-3725-8682](https://orcid.org/0009-0009-3725-8682)

Revised August 2026

Central claim


Complexity by itself does not identify a designer. Artifacts support agency because specified agent hypotheses predict discriminating traces within an appropriate comparison class, while caves and evolved systems can be complex under known non-agent mechanisms. A responsible inference must therefore compare causal pathways and let trace evidence—not wonder—move the odds. This blocks the complexity-to-designer shortcut; it does not rule out every design hypothesis or every disciplined agency inference.

Preview:

Starting point	Hidden bridge	Comparative pressure	Rational conclusion
Some complex artifacts are designed.	Complexity is assumed to remain diagnostic when the target class changes.	Natural mechanisms also generate intricate structure, while genuine artifacts carry additional production traces.	Complexity can prompt inquiry, but agency requires a specified rival comparison and discriminating evidence.

The result targets a shortcut from complexity to agency. It leaves open design arguments that supply independent specifications, comparative probabilities, and risky trace predictions.

Visual preview:



The image moves from cave complexity to a natural production process, contrasts that case with a known artifact class, and depicts the unsupported leap back to design. It isolates the transfer error; it does not claim that every agency inference fails.

Abstract

Complexity alone cannot warrant a designer. The familiar move from watches, books, machines, or codes to caves, organisms, and cosmic structure changes the comparison class while preserving a correlation established inside an artifact-rich sample. That transfer is invalid as a universal deduction and evidentially weak unless the target feature is more expected under a specified agent hypothesis than under specified non-agent alternatives. Known cave formation, evolutionary accumulation, co-option, and constrained material processes show why intricacy is not exclusive to agency. They do not prove that no designer exists; they raise the probability of complex outcomes under live non-agent pathways and thereby remove the shortcut.

The paper reconstructs the cave analogy around four requirements: representative comparison classes, explicit causal alternatives, discriminating traces, and exposure to contrary evidence. It then tests stronger proposals. Specified complexity and irreducible complexity improve on mere wonder because they attempt to define a narrower target, but they still require justified probability models, independent specifications, and serious comparison with evolutionary pathways. Fine-tuning arguments are a different family and are not refuted by the cave analogy, though they face their own comparative burdens. Archaeology and forensics show what successful agency inference looks like: investigators use contextual traces, production sequences, causal access, and rival models rather than complexity alone. The conclusion is methodological and bounded. Complexity may identify a question worth investigating; only comparative trace evidence can identify agency.

Keywords: design inference; complexity; reference class; likelihood; intelligent design; evolution; artifact detection

1. The Cave Case and Its Exact Target

Imagine a father and child entering a cavern. The father has seen mines, tunnels, and carved chambers. He points to branching passages, vaulted spaces, repeated patterns, and functional routes, then concludes that caves are designed. The child notices water channels, mineral deposits, irregular fractures, sediment, and the continuity between the walls and the surrounding rock. Some underground spaces are excavated, the child grants, but natural processes also produce caves.

The story is memorable because the father's premise contains a truth. Human beings do design tunnels. We can recognize many of them reliably. The mistake lies in converting a correlation inside a familiar subclass into a rule for a broader class whose production is disputed. The inference has roughly this shape:

1. Many familiar complex structures—watches, books, machines, buildings, and tunnels—are designed.
2. The newly encountered target is also complex.
3. Therefore the target is designed.

The conclusion does not follow from those premises. The designed examples were selected from classes in which agency was already established by independent history and traces. Their complexity accompanies design; it did not by itself establish design. When the target changes from watches to cells or from mines to limestone caverns, the old correlation cannot simply be carried over.

This is the *cave-design error*: transferring an agency association from a design-saturated reference class to a disputed natural class while treating the shared word “complex” as the entire bridge. The title-page visual compresses that sequence. Cave complexity is set beside natural formation, a watch represents a known artifact class, and the final panel marks the leap that still requires an argument.

The target must remain exact. The paper does not argue that agency can never be inferred, that natural mechanisms exclude a creator, or that every teleological argument reduces to crude complexity. Paley's watch analogy is historically important because it directs attention to organized functional relations, while Hume's earlier critique already pressed the weakness of moving from a limited human comparison to the order of nature (Hume, 1998; Paley, 1802). Contemporary arguments involving fine-tuning, specified patterns, or system-level

dependency add claims not contained in the bare premise “this is complex”. Those additions must be assessed on their own merits.

The limited conclusion

One naturally formed complex object refutes the universal rule that complexity entails design. It does not refute every particular design hypothesis. The stronger methodological result is that complexity alone is not a discriminator; any successful agency inference must obtain its force from additional comparative evidence.

2. Two Failures Often Blurred Together

2.1 The invalid universal transfer

The first failure is deductive. From “all complex objects in a selected artifact sample are designed” it does not follow that all complex objects are designed. The sample condition has been dropped. A naturally formed cave, snow crystal, river network, weather system, or evolved organ is enough to defeat the universal generalization. Darwin’s achievement was not to deny biological intricacy but to propose a process that can accumulate adaptive structure without foresight (Darwin, 1859).

This point is decisive against slogans such as “complexity requires a designer”. If even one non-agent process can generate a complex outcome, complexity is not a design-exclusive property. Redefining the target may produce a new argument, but it also concedes that ordinary complexity did not carry the original one.

2.2 The weak probabilistic transfer

The second failure is evidential and needs more care. A property need not be exclusive to one cause before it counts as evidence. Smoke can support fire even though fog and machines also produce visible plumes. What matters is comparative expectation. Let E be the observed feature, D a specified design hypothesis, and N a specified non-agent hypothesis. The feature favors D over N only when:

$$\frac{P(E \mid D)}{P(E \mid N)} > 1.$$

In ordinary language, the evidence must be more expected if the design hypothesis is true than if the rival is true. The size of that difference determines how much the observation moves the comparison. Likelihood reasoning of this kind is central to philosophical analysis of evidence in evolution and design (Sober, 2008).

Merely calling E complex supplies neither side of the comparison. A vague designer with no specified aims, constraints, materials, or production method does not determine $P(E \mid D)$. Ignoring natural mechanisms does not determine $P(E \mid N)$. The word “complex” names the response to the object, not the probability model that makes it evidence.

This is where reference classes matter without doing all the work. The watch sample has a high design rate because it was assembled from known artifacts. That rate cannot be transferred to organisms or universes whose classification is the very question. But even a representative reference class would not settle a particular case without relevant features. The investigator still needs to ask what the competing causes predict in the target domain.

Table 1: Where the cave-design transfer fails.

Stage	What is actually known	What is silently added	Correct question
Artifact sample	Watches and tunnels have independent production histories.	Complexity is treated as the feature that established those histories.	Which traces, contexts, and sequences identified manufacture?
Target transfer	A cave, organism, or system is intricate.	The artifact sample is treated as representative of the new domain.	Which comparison class and causal pathways fit this target?
Universal conclusion	Some complex things are designed.	Every complex thing is assumed to be designed.	Is complexity actually exclusive to agency?
Evidential conclusion	Complexity can accompany agency.	Complexity is assumed to be much more likely under design.	How probable is this feature under each specified rival?

The cave analogy is therefore more than a counterexample and less than a disproof of design. It exposes a missing bridge. Universal arguments fail because the bridge is false. Probabilistic arguments remain incomplete until the bridge is quantified or otherwise supported by a real causal comparison.

3. Natural Mechanisms Change the Evidential Meaning of Complexity

Complexity can feel like a positive trace of agency when the alternative is pictured as instantaneous disorder. That contrast is misleading. Non-agent alternatives are not one hypothesis called “chance”. They include lawful self-organization, cumulative selection, inheritance, constraint, drift, co-option, feedback, and historical path dependence. Different mechanisms apply to different targets.

Karst caves provide the cleanest illustration. Slightly acidic water moves through fractures and bedding planes in soluble rock, enlarging openings and reorganizing drainage. Over time, hydrology, chemistry, collapse, sediment transport, and mineral deposition produce extensive passage networks. U.S. Geological Survey work documents dissolution as the primary process creating karst landscapes and their caves (Kappel et al., 2020). Once that pathway is known, complexity no longer separates excavation from geology. Tool marks, blast holes, spoil, lining, plans, and historical access may still distinguish a tunnel; passage count and visual impressiveness do not.

Biological complexity requires different mechanisms. Cumulative evolution is not a single random draw at a finished system. Variation occurs within populations; inheritance preserves differences; selection, drift, constraint, duplication, and co-option alter lineages over time. Digital-evolution experiments have shown complex logical functions arising by building on earlier functions, with historical stepping-stones affecting later outcomes (Lenski et al., 2003). Digital organisms are human-built models, so they are not direct demonstrations of every biological pathway. They establish the narrower point that cumulative selection can assemble functions that a one-step improbability calculation would badly misdescribe.

Biology supplies concrete pathway work as well. Co-option gives existing genes, structures, or regulatory circuits new roles (True & Carroll, 2002). Comparative research on bacterial flagella identifies homologous components, related secretion systems, and multiple evolutionary questions that replace the picture of an all-or-nothing machine (Pallen & Matzke, 2006). Such work may leave historical details unsettled. Its evidential significance is that the non-agent likelihood is not responsibly modeled as the chance assembly of a present system from unrelated parts.

Table 2: Mechanisms that alter the comparison without proving more than they show.

Target	Non-agent pathway	What the evidence contributes	What it does not establish
Karst cave	Dissolution, groundwater flow, fracture enlargement, collapse, and deposition.	Complex passages are expected without excavation.	No cave was ever modified or built by agents.
Complex function in a model	Reproduction, mutation, competition, retention, and cumulative selection.	Multi-step functions need not appear in one random leap.	Every biological novelty followed the model's exact path.
Molecular system	Duplication, co-option, divergence, homologous components, and historical constraint.	Present dependency does not by itself show evolutionary inaccessibility.	Every transition has already been reconstructed in complete detail.
Stone form	Fracture mechanics, reduction sequence, context, and constrained production space.	Visual complexity may arise without the degree of foresight an observer projects onto it.	Human manufacture can never be inferred from lithic evidence.

The last row matters because even genuine artifacts must be interpreted cautiously. Experimental archaeology shows that apparently sophisticated stone forms can emerge under severe constraints without the elaborate advance planning that modern observers may attribute to them (Moore, 2023). Agency may be secure while the inferred kind or depth of agency remains overstated. The lesson is general: an effect supports only the causal features that make it comparatively expected.

Natural mechanisms therefore do not function as a philosophical veto on design. They change the evidential landscape. As $P(E \mid N)$ rises, the same feature supplies less support for D over N . A design advocate may respond with a narrower feature or a stronger model, but the response must engage the mechanism rather than return to the original feeling of improbability.

4. How Genuine Artifact Inference Works

Archaeologists, forensic investigators, engineers, and historians routinely infer agency. Their success is not an embarrassment to the cave critique; it reveals the missing standard. They do not normally classify an item as manufactured because it is intricate in the abstract. They combine several kinds of evidence:

- a plausible agent with the capacity and causal access to produce the target;
- a production sequence that connects actions to the observed result;
- traces that differ under manufacture and under live non-agent rivals;
- context, such as associated tools, records, waste, repeated conventions, or location;
- and exposure to failure, including observations that would lower confidence in agency.

A broken stone may be simple yet strongly support manufacture if its flake scars, platform, use wear, context, and reduction sequence fit human work better than accidental fracture. A crystal may be visually elaborate yet strongly support a non-agent process because its lattice structure and growth conditions fit crystallization. Complexity and design evidence can therefore move in opposite directions.

Table 3: A trace-based agency standard.

Question	Why it matters	Warning sign
Which agent hypothesis?	Capacity, aims, constraints, and access determine what agency predicts.	A designer able and willing to produce any outcome predicts too little.
Which non-agent rivals?	Evidence is comparative; omitting a live mechanism inflates the apparent need for agency.	The rival is caricatured as a single random assembly.
Which traces discriminate?	Independent production marks can favor one pathway over another.	The only reported trace is complexity itself.
Was the criterion fixed in advance?	Prior specification reduces the risk of treating any surprising pattern as intended.	The target property changes after each mechanism is supplied.
What would count against agency?	A hypothesis that risks loss can gain public evidential force.	Every outcome is redescribed as compatible with hidden intentions.

The trace-over-awe rule

Do not ask first whether an object is too impressive to be natural. Ask which specified process predicts its actual traces better than the live alternatives do, and identify in advance what evidence would reverse the comparison.

This standard also explains why a vague divine designer is difficult to use in likelihood reasoning. A *God-concept* will mean here a specified conception of a deity, including the relevant aims, capacities, constraints, and relation to the target. Theism in general does not supply one probability distribution over organisms, caves, or physical constants. Different God-concepts may predict elegance, plenitude, hiddenness, autonomy, suffering, intervention, regularity, or combinations of these. Unless the operative conception and bridge principles are stated, almost any observation may be accommodated after the fact.

Accommodation is not prediction. A hypothesis can remain logically compatible with a feature without making that feature more expected than its rivals do. The same distinction applies outside theology. An unknown human builder could have wanted an irregular water-carved passage, but that possibility does not make human excavation the better explanation when geology already predicts the traces.

5. Stronger Design Proposals

Sophisticated design advocates rarely rest with the sentence “this is complex”. Their refinements deserve to be stated in their strongest relevant form because each one changes the argument.

5.1 Specified complexity

Dembski’s design inference combines low probability under relevant non-agent hypotheses with conformity to an independently given pattern (Dembski, 1998). This is a genuine improvement over ordinary complexity. A long random string is improbable but not thereby designed; a pattern fixed independently of the outcome may be more evidentially interesting.

The improvement also reveals the remaining burdens. The probability must be computed under an adequate rival model rather than a one-step chance caricature. The specification must be independent enough to avoid choosing a pattern after seeing the result. The set of relevant non-agent hypotheses must include cumulative and

structured processes. Finally, eliminating one rival does not automatically identify a particular agent. Those are model questions, not semantic objections to the word “specified”.

5.2 Irreducible complexity

Behe’s irreducible-complexity proposal focuses on systems whose present function depends on several interacting parts (Behe, 1996). It too abandons bare complexity for a more specific dependency claim. But present-day indispensability is not the same as historical inaccessibility. Parts may have had earlier functions; systems can be built through duplication, co-option, scaffolding, redundancy, and later loss. The relevant test compares evolutionary pathways and evidence, not a mental attempt to remove one current part.

Research on co-option and flagellar evolution directly addresses that distinction (Pallen & Matzke, 2006; True & Carroll, 2002). Philosophical criticism has also documented shifts between definitions of irreducibility and the risk of turning the target into a moving boundary (Boudry et al., 2010). None of this licenses saying that every molecular history is complete. It licenses requiring the defender to show that the relevant pathways are not merely unknown but genuinely implausible under a properly specified evolutionary model.

5.3 Fine-tuning and a different target

Cosmological fine-tuning arguments concern the apparent life-permitting range of physical parameters, not the ordinary inference from visible complexity. The cave analogy does not refute them. A fine-tuning case may still need a justified probability measure, a specified design hypothesis, treatment of selection effects, and comparison with rival cosmological models, but those are additional debates (Sober, 2008). Keeping them separate prevents a victory or failure in one family from being smuggled into another.

Table 4: What stronger design proposals add and still owe.

Proposal	Genuine improvement	Remaining burden	Reach of the cave critique
Specified complexity	Adds prior pattern and probability, not complexity alone.	Defensible rival model, independent specification, and agent identification.	Defeats only a return to wonder or an artifact-biased sample.
Irreducible complexity	Names present functional dependency among parts.	Evidence that historical co-option, redundancy, and pathway change are implausible.	Shows that present intricacy is not itself historical impossibility.
Fine-tuning	Targets parameter ranges rather than visible intricacy.	Probability measure, observation selection, rival cosmologies, and specified aims.	Does not directly refute the argument.
Trace-based agency	Predicts contextual production marks and risks contrary evidence.	Reliable traces, causal access, and serious alternatives.	Satisfies rather than violates the paper’s methodological demand.

Refinement is not automatically goalpost movement. Lakatos distinguishes research programmes that generate new predictions and discoveries from those that protect a core claim mainly through after-the-fact adjustments (Lakatos, 1970). A design refinement is progressive to the extent that it fixes criteria before inspection, exposes the view to contrary evidence, and discovers traces that rivals do not predict. It is degenerating when “designed complexity” becomes whatever remains unexplained today.

6. A Worked Comparative Protocol

The cave analogy becomes useful when it supplies a procedure rather than a slogan. A disciplined comparison can proceed in five steps.

1. **Define the target feature.** Replace “complex” with observable structure: branching geometry, tool-like marks, functional integration, sequence regularity, or another exact property.
2. **Specify the causal rivals.** Name the agent hypothesis and the strongest relevant non-agent processes. Do not compare design with an empty label such as chance.
3. **Choose discriminating traces before concluding.** Ask what each pathway should leave behind and which traces are independent of the original classification.
4. **Compare expectations.** Determine whether the observed traces are appreciably more expected under one rival. Qualitative comparisons can be useful, but they must still identify direction and reason.
5. **State reversers and limits.** Say what evidence would reduce confidence, which rivals remain live, and how narrow the conclusion is.

Consider a straight underground passage. Under a human-excavation hypothesis, one might expect cut or blast marks, standardized dimensions, support structures, spoil, access routes, records, or alignment with a human project. Under a karst hypothesis, one might expect lithological continuity, dissolution features, drainage alignment, mineral deposition, fractures, and sediment patterns. Straightness alone may occur under both. A collection of production traces can move the comparison.

Now consider a molecular machine. “It has many coordinated parts” is not yet a pathway comparison. A design hypothesis must say why the relevant God-concept or other agent would produce those parts, in that arrangement, with that history and distribution. Evolutionary hypotheses must identify homologues, precursor functions, co-option routes, population conditions, constraints, and expected imperfections. A historical explanation may remain incomplete on both sides. The rational outcome may then be suspended judgment or a graded likelihood comparison, not an automatic verdict.

This protocol also disciplines negative arguments. Failure to produce a complete natural history can support design only when the search was capable of finding the relevant pathway, the alternatives are sufficiently exhaustive, and the design hypothesis positively predicts the observed traces. Otherwise absence of explanation is mainly information about the present state of inquiry.

The standard is demanding because agency inference is causal inference. Causal explanations earn their status by connecting a process to a pattern and outperforming alternatives. Complexity can be part of the pattern described. It cannot substitute for that work.

7. Objections and Replies

7.1 Natural processes may themselves be designed

A theist may hold that evolution, physical law, and cave-forming geology are instruments of a creator. Nothing in this paper disproves that view. The point is inferential: once the same complex outcome is expected through a natural process whether or not an ultimate designer established the process, complexity no longer distinguishes the ultimate hypotheses. The design claim has moved from a local artifact inference to a broader metaphysical explanation and needs evidence appropriate to that level.

7.2 A counterexample does not show that complexity is never evidence

Correct. A naturally complex object refutes a universal entailment, not every probabilistic relation. Complexity could favor design in a carefully specified comparison. But that is exactly why the likelihoods, reference class, agent aims, and non-agent pathways must be supplied. The cave error is treating complexity as self-interpreting evidence.

7.3 Natural mechanisms do not explain everything

Also correct. Evolutionary and geological explanations have open questions, and particular histories can be difficult to reconstruct. An explanatory gap is not nothing, but its direction is underdetermined until the alternatives make contrasting predictions. “No complete pathway is known” does not by itself mean “agency is probable”, especially when an unspecified designer predicts no unique trace.

7.4 Archaeology and forensics infer design without observing production

They often do, and they illustrate the paper’s positive standard. Investigators use independently understood agents, production constraints, contextual records, trace regularities, and comparison with non-agent causes. Direct observation of manufacture is unnecessary when the causal model predicts discriminating remains. Complexity alone is unnecessary as well.

7.5 A designer need not reveal motives or mechanism

Full knowledge is not required. Ordinary inference works under uncertainty. But less specification reduces evidential resolution. If an agent could have any motive, use any method, and produce any outcome, then the hypothesis can accommodate everything and comparatively predict little. A modest, testable design hypothesis may be stronger evidence than a grand but unconstrained one.

7.6 The cave analogy is stacked because cave formation is already known

That feature is pedagogically useful, not illicit. The story lets readers see how mechanism changes the meaning of complexity in a settled case. The analogy does not assume that every biological or cosmological mechanism is equally settled. It transfers a rule of inquiry: before moving from intricacy to agency, compare causal pathways and traces. The confidence appropriate to each target can then differ with the evidence.

8. Conclusion

The cave-design error begins with a truth and ends with an unsupported transfer. Some complex things are artifacts. Yet their complexity did not independently establish their manufacture, and their design-rich class does not represent every natural target. Once natural mechanisms make complex outcomes expected, wonder loses its power to discriminate.

The rational repair is not a ban on agency. It is a higher standard for inferring it: specify the agent and non-agent rivals, identify traces that differ between them, compare expectations, and state what would count against the conclusion. Stronger design proposals deserve separate assessment precisely because they add claims beyond ordinary complexity.

Complexity is an invitation to investigate. It is not a license to conclude. Readers should therefore ask one question whenever the cave-design move appears: what evidence points to this causal pathway rather than merely describing the object as impressive?

A. Formal Reconstruction of the Central Thesis

Premise 1. If complexity alone warrants an inference to agency in a target domain, then complexity must discriminate agent production from the live non-agent processes in that domain.

Premise 2. Complexity does not provide that discrimination. Known non-agent processes generate complex outcomes, and the familiar artifact examples come from classes in which design was established through additional history and traces.

Premise 3. An agency inference gains evidential force only when a specified agent hypothesis predicts the observed traces better than specified rivals and remains exposed to contrary evidence.

Conclusion. Therefore complexity alone does not warrant a designer. Any successful design inference must receive its force from additional comparative trace evidence.

Validity. The conclusion follows from the three premises as a necessary-condition argument. Premise 1 states what would be required for complexity alone to carry the inference; Premise 2 denies that condition; Premise 3 identifies what must replace it. The argument does not conclude that agency is absent.

Support for the premises. Premise 1 is conceptual: evidence must distinguish the hypotheses it is used to compare. Premise 2 is supported by cave geology, evolutionary mechanisms, and the independent traces used to recognize artifacts. Premise 3 is methodological and receives support from likelihood-based philosophy of evidence and successful trace-based agency inquiry. These premises strongly constrain the shortcut, but they leave particular design hypotheses open to evaluation when those hypotheses supply the missing comparison.

References

- Behe, M. J. (1996). *Darwin's black box: The biochemical challenge to evolution*. Free Press.
- Boudry, M., Blancke, S., & Braeckman, J. (2010). Irreducible incoherence and intelligent design: A look into the conceptual toolbox of a pseudoscience. *The Quarterly Review of Biology*, 85(4), 473–482. <https://doi.org/10.1086/656904>
- Darwin, C. (1859). *On the origin of species by means of natural selection*. John Murray.
- Dembski, W. A. (1998). *The design inference: Eliminating chance through small probabilities*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511570643>
- Hume, D. (1998). *Dialogues concerning natural religion* (R. H. Popkin, Ed.; 2nd ed.). Hackett Publishing Company.
- Kappel, W. M., Reddy, J. E., & Root, J. C. (2020). *Statewide assessment of karst aquifers in new york with an inventory of closed-depression and focused-recharge features* (Scientific Investigations Report No. 2020-5030). U.S. Geological Survey. <https://doi.org/10.3133/sir20205030>
- Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge* (pp. 91–196). Cambridge University Press. <https://doi.org/10.1017/CBO9781139171434.009>
- Lenski, R. E., Ofria, C., Pennock, R. T., & Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423, 139–144. <https://doi.org/10.1038/nature01568>
- Moore, M. W. (2023). Hominin stone tool manufacture: Design space constraints and implications for interpreting cognitive evolution. In T. Wynn, K. A. Overmann, & F. L. Coolidge (Eds.), *The oxford handbook of cognitive archaeology* (pp. 675–690). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780192895950.013.20>
- Paley, W. (1802). *Natural theology; or, evidences of the existence and attributes of the deity, collected from the appearances of nature*. R. Faulder.
- Pallen, M. J., & Matzke, N. J. (2006). From the origin of species to the origin of bacterial flagella. *Nature Reviews Microbiology*, 4(10), 784–790. <https://doi.org/10.1038/nrmicro1493>
- Sober, E. (2008). *Evidence and evolution: The logic behind the science*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511806285>
- True, J. R., & Carroll, S. B. (2002). Gene co-option in physiological and morphological evolution. *Annual Review of Cell and Developmental Biology*, 18, 53–80. <https://doi.org/10.1146/annurev.cellbio.18.020402.140619>