

Aaron Berney

PHIL 430

Professor James Woodbridge

March 15, 2026

Why Confirmation Might Be Dead

American philosopher Nelson Goodman posited that in order to develop an adequate confirmation theory, one established on rational, logically sound, and generally applicable grounds, we must first address the long-standing problem of induction prompted by David Hume. To explain what Goodman means, we must first provide the necessary context and background surrounding the philosophical discourse of confirmation theory. To the uninitiated, science may seem straightforward: propose a hypothesis, develop an experiment to test the hypothesis, collect the empirical results, and if the results mirror the prediction, we accept the hypothesis. However, if we brush aside the procedural surface of scientific practice and examine its epistemological foundation, we will note that it becomes significantly more complex. When we question what confirmation is on a philosophical basis, what it means for evidence to actually constitute support for a hypothesis, we find that the answer seems to be indeterminate.

This uncertainty surrounding confirmation stems from its complete reliance on inductive reasoning. In contrast to deductive reasoning, which operates *a priori* and provides conclusions that follow directly from their premises (like how the statement “all bachelors are unmarried” is true by definition), induction operates differently. Inductive reasoning operates on reasoning frameworks like inference to the best explanation and Occam’s Razor rooted in past observation. Inductive reasoning reduces to making inferences about the future based on what we have previously observed to be true. And this means that every scientific theory developed outside the formal sciences, every empirical conclusion we have made, is not logically certain. This has an unsettling implication. There is no guarantee that the sun will rise tomorrow. There is no guarantee that the chair I am sitting on will not disappear in the next

second. We can raise any number of absurd statements, and yet we cannot doubt them with logical certainty. Habit and repetition have molded our intuition.

Inductive reasoning lacks logical certainty as a consequence of its circular nature. This is the problem given by philosopher David Hume: in order to justify using the past as evidence for the future, we must already assume that the past is a reliable source from which to draw inferences about the future. In other words, we can only justify induction through an inductive assumption. This is what Goodman referred to when he mentioned the need to address the problem of induction raised by Hume: in order to develop an adequate theory of confirmation, we must logically justify the use of inductive reasoning despite its circularity. However, a response to Hume's problem emerged: deductive reasoning faces the same issue with circularity as inductive reasoning. It too must rely on deductive principles to justify its own structure. If philosophers permit the use of deduction despite circularity, they must equally permit the use of induction. Goodman took this to imply that Hume's problem was not as much of a challenge to scientific activity as some had thought, and he proceeded to develop an even more aporetic riddle of induction. But before arriving at his riddle, we must examine two preliminary problems, each one exposing a distinct failure in the LE framework, and together setting the stage for the most formidable opponent of confirmation theory: Goodman's riddle.

With Hume's circularity problem addressed, new questions arose. The question became how to distinguish good inductive inferences from bad ones, and what made one piece of evidence more supportive of a hypothesis than another. These questions gave rise to several new puzzles such as Carl Hempel's "Ravens Paradox" and the tacking problem. Consider the hypothesis "all ravens are black." We can naturally assume that every black raven we observe adds support to this hypothesis. But if we rephrase the hypothesis, we get the equivalent hypothesis "all non-black things are non-ravens." And this means any observation of a non-black item that is not a raven counts as supportive evidence for the hypothesis that all ravens are black. A white shoe, yellow taxi, blue t-shirt, etc. all technically count as

evidence. While the logical structure holds up, it seems absurd to conclude that observing white shoes confirms the raven hypothesis.

The tacking problem reveals another issue: when we confirm a hypothesis through empirical observations, we never just confirm one isolated claim. Every hypothesis unavoidably carries an infinite amount of sub-hypotheses that exist as background assumptions. If we hypothesize that the sun will rise tomorrow and it does, that observation also confirms the conjunction of that hypothesis with any number of unrelated but simultaneously true claims: that a particular grain of sand on a Hawaiian beach was in its precise location, that a particular fish is located in a specific body of water, that a philosophy professor teaching PHIL 430 at UNLV is white, etc. The confirmatory power of evidence bleeds indiscriminately into everything tacked onto the original hypothesis. Observation turns out to be a remarkably imprecise instrument.

With those two problems contextualized, we can introduce the most vexing problem: Goodman's riddle. Goodman wondered what made one inductive hypothesis more projectible than another. To demonstrate his issue, he developed the infamous predicate "grue." Goodman defined grue as being green and examined before some time in the future t , or blue if examined after time t . Next, consider the hypothesis "all emeralds are green" and "all emeralds are grue." Given t = the year 2050, it seems to deviate from natural thought to assume that emeralds would suddenly appear as blue instead of green after the year 2050 as given in the grue hypothesis. However, every single green emerald observed now supports both hypotheses equally. The evidence and logical structure are identical, yet it seems deceptively unreasonable to believe in the grue hypothesis. There is no way to logically demonstrate why "green" is a better predicate than "grue." Both are formally valid and consistent with available evidence.

With no way to rationally favor one projection over another, the task of developing an adequate confirmation theory lies in significant peril. To put into perspective just how damaging Goodman's riddle was, we must better define what Logical Empiricists sought to develop with an "adequate" theory of confirmation. They had a clear goal: to develop a theory of confirmation grounded in rationality, logic,

and generalizable across all scientific domains. They wanted confirmation theory to operate under a formal, systematic, and syntactic framework that could definitively explain when and how evidence supports a hypothesis. They sought a precise, rigorous logical structure that could unify scientific reasoning across the board.

The aforementioned problems make it painfully obvious why their endeavor proved to be so hard to attain. The Ravens Paradox showed that even the most straightforward logical equivalences generate absurd confirmatory relationships. The tacking problem showed that confirmation bleeds indiscriminately across conjunctions. And Goodman's new riddle of induction showed that we cannot even determine which predicates are legitimate candidates for inductive projection. Faced with the issues his own riddle introduced, Goodman developed what we can only describe as an explanation; as will become clear, explanation is precisely the correct word, and solution is not. His explanation centered on the concept of "entrenchment". He argued that certain predicates are more entrenched than others. He defined an entrenched predicate as one that speakers have widely and repeatedly used across the history of human language and thought. In other words, these are concepts that are so normalized and embedded in our understanding of the world that they serve as natural features of reality. In the context of his riddle, "green" would be entrenched and "grue" would not be, which is why we instinctively select the green hypothesis over the grue hypothesis. Goodman's response surrendered the answer to linguistics and sociological phenomena, not logic.

To illustrate how contingent his response really was, Goodman described a hypothetical society that had grue as the entrenched concept instead of green. This was a society that had always described objects in reference to their future color properties relative to a specific point in time. In this society, the grue hypothesis would seem like the obvious choice, and green would seem like the foreign predicate. Green would be just as puzzling and arbitrary as we find grue in our own society. Entrenchment is a product of human convention, usage, and history. There is no relation to logic and rationality.

This is exactly why Goodman's response cannot deal with the problems his own riddle introduced. It provides an explanation for why we feel the way we do, but it does not tell us whether or not we are justified in feeling that way on any logical grounds. What emerges from all of this is an uncomfortable picture: confirmation theory, as the Logical Empiricists envisioned it, appears to be inherently irrational. There is no logical principle that cleanly separates good inductions from bad ones, no formal rule that tells us which evidence genuinely supports which hypothesis, and no rational basis for preferring green over grue. Goodman's New Riddle of Induction deals a devastating blow to the idea of confirmation that the Logical Empiricists pursued.

However, even if confirmation may not be rationally justifiable in the strict logical sense, confirmation remains a pragmatically indispensable tool in science. This makes the pursuit of an adequate confirmation theory even more imperative. Perhaps confirmation theory is not finished with Goodman's Riddle. It may simply mean that the angle of attack needs to change entirely. Rather than search for a purely logical, syntactic theory of confirmation, perhaps what science needs is a probabilistic one. This is precisely where Bayesian confirmation theory emerges. Bayesianism takes a new outlook on confirmation as a matter of degrees of belief that update over time. It treats the probability of a hypothesis as a value that updates continuously as new evidence accumulates. The way it operates is through Bayes' Theorem: an equation that calculates the probability of a hypothesis based on given evidence. It is given by $P(H|E) = P(E|H) \cdot P(H) / P(E)$. $P(H|E)$ describes the probability of hypothesis H being true given the observed evidence E. $P(E|H)$ describes the probability of observing evidence E if hypothesis H were to be true. $P(H)$ describes the probability of hypothesis H being true before observing evidence E. Finally, $P(E)$ describes the probability of observing evidence E regardless of whether or not hypothesis H is true.

The equation allows for the probability of a given hypothesis to update continuously as new evidence is discovered. Additionally, the percent change in a hypothesis's probability before and after a new piece of evidence is provided also gives weighted values between different pieces of evidence. Previous approaches to confirmation theory treated it as a binary logical relationship between evidence

and hypotheses that inherently made the justification of confirmation on any logical and rational grounds impossible. Rather than asking whether evidence confirms a hypothesis in an absolute logical sense, Bayesianism questions how much a piece of evidence should shift our confidence in it. This provides a far more flexible and honest picture of how scientific reasoning actually works. Evidence provides belief on a spectrum, not in absolute terms.

This gives Bayesianism a distinct advantage over the idea of confirmation theory sought by the Logical Empiricists. It accommodates the Raven's Paradox by giving a way to evaluate the supportiveness of evidence, and it provides a more accurate account of how scientists actually evaluate scientific theories. However, Bayesianism also carries its own flaws that serve as insurmountable obstacles to confirmation theory. The entire framework Bayesianism operates upon is fundamentally subjective. The equation operates on prior probabilities to calculate the degree of belief we should have in a hypothesis, and there is no objective, universally agreed-upon method for determining what a prior should be. Two scientists examining the same hypothesis can assign entirely different prior probabilities based on their individual backgrounds, intuitions, and theoretical commitments, and no finite amount of evidence can negate the subjectivity of a prior.

Bayesians respond to this objection with the concept of convergence. If given enough evidence over time, Bayesians argue that scientists who have different priors for a hypothesis will gradually update their beliefs to reflect the same priors and converge on the truth (the true, objective priors) as more and more evidence accumulates. However, this explanation is even poorer than Goodman's answer to his own riddle. This idea of convergence is unreachable. It takes place on an infinitely long timeline. The time it takes to collect the amount of evidence needed for true convergence on just a single prior in a single hypothesis is infinite. Science needs to progress now. It cannot wait for convergence in the distant future. But the real obstacle is not the impracticality of waiting for convergence. It is once again Goodman's riddle. Bayesianism does not solve Goodman's riddle, but it reframes the problem. Faced with the problem of grue, a Bayesian may argue that it is not obviously derailed by Goodman's riddle in the same

way as the Logical Empiricists' binary perspective was. The grue hypothesis, under a Bayesian framework, would simply receive a very low prior probability because nothing in our accumulated experience gives us reason to assign it serious credence. However, this only restates the riddle: if Bayes' Theorem rests on the subjective priors of human scientists, then the only reason grue gets a low prior is because of entrenchment and conventionality. The only thing that Bayes' Theorem achieves in the end is to provide a formalization of subjectivity and reimagine Goodman's riddle. Goodman's riddle sinks the idea of rational, logical confirmation from both a binary and probabilistic perspective.

This paper reaches an uncomfortable conclusion. It implies that all scientific endeavors are grounded in irrationality. Hume showed us that induction cannot justify itself without circularity. Hempel's Ravens Paradox showed that logical confirmation generates absurd results from valid steps. The tacking problem showed that evidence confirms far more than we intend. And Goodman's New Riddle showed that we cannot even determine which predicates are legitimate candidates for inductive projection in the first place. Taken together, these problems do not just challenge confirmation theory; they suggest that confirmation theory may be impossible to justify on any valid epistemological grounds. And yet, science persists. Scientists form, test, and either trust or discard hypotheses, and this process has proven extraordinarily productive, regardless of its philosophical uncertainty. Perhaps this suggests that the problem is not with confirmation as a practice, but with the framework we use to evaluate it. Instead of asking what makes evidence confirm hypotheses, it may be more philosophically productive to ask what makes evidence disconfirm hypotheses. This is the core principle behind philosopher Karl Popper's falsificationism. If we cannot confirm hypotheses, we can, in principle, still refute them. However, whether falsification succeeds where confirmation has failed is a question beyond the scope of this essay. But it may be the most promising direction left to explore. Perhaps science can still find sound epistemological roots as a philosophically rational and logical practice.