Ravens, Shoes and the Structure of Confirmation

John Quiggin

December 2025

Abstract

Hempel's raven paradox arises only under assumptions that abstract away from the structure of finite domains of inquiry. Once that structure is restored, the paradox dissolves. Observing a non-raven confirms a universal generalisation only if the prior distribution embeds a substantive empirical relationship between ravens and non-ravens. Without such a relationship, no confirmation occurs. The finite-case models developed here make this requirement transparent and connect with earlier insights due to Hosiasson-Lindenbaum, Good, and Vranas. The broader lesson, reinforced by examples from natural science and econometrics, is that confirmation is always model-relative.

1 Introduction

Hempel's raven paradox [Hempel, 1945] remains one of the most widely discussed puzzles in the theory of confirmation. The paradox arises from the logical equivalence of:

(H) All ravens are black

and

(H*) All non-black things are non-ravens.

If a black raven confirms (H), then a non-black non-raven should confirm (H*), and hence (H). Bayesian analyses typically treat this as an unavoidable result [Maher, 1999, 2004].

This paper argues that the apparent paradox arises only because confirmation is analysed using abstract priors and universal quantifiers, which obscure the implicit assumptions about empirical structure. Once that structure is made explicit, the paradox disappears.

2 Two examples

Two simple examples illustrate the point that the validity or otherwise of Hempel's claim depends on assumptions about empirical structure. In simple finite cases, the fact that some such assumptions must be made is obvious.

2.1 Example 1

Consider a world with one raven and one shoe. Each may be black or non-black. If the colour of the shoe is independent of the colour of the raven, observing the shoe tells us nothing about the colour of the raven. Logical equivalence alone does not generate confirmation—precisely the point anticipated in the early relevance condition of Hosiasson-Lindenbaum [1940].

This extends to any finite numbers of ravens and shoes. If colours are independent across categories, no observation of a shoe affects the probability that all ravens are black. Moreover, we may assign any distribution over the number of ravens and shoes; since independence blocks confirmation in every finite configuration, it blocks confirmation under any mixture over them.

2.2 Example 2

There is one finite setting in which observing a non-raven does confirm the universal: sampling without replacement. Suppose two cards are drawn from a standard 52-card deck, one labelled the circle card (non-raven), one the star card (raven). If the circle card is red, the probability that the star is black increases from 1/2 to 26/51.

This example instantiates the condition identified by Vranas [2004]. He shows that non-black non-ravens confirm the universal only if the prior makes worlds in which all ravens are black systematically associated with worlds in which non-ravens are less black. The hypergeometric structure of drawing without replacement generates exactly this correlation.

2.3 Good's Example and Hidden Structure

Good's "white shoe" example [Good, 1966] illustrates the same point. His urn model shows that non-black non-ravens confirm the universal only under priors that correlate the colour patterns of ravens and non-ravens. In this sense, Good's construction has exactly the same structure as the finite models developed here.

But Good's approach is sufficiently elaborate—urns, proportions, hierarchical priors—that it was easy for Hempel [Hempel, 1945] to dismiss it as "building in" the very structure that generates the confirmation. The finite models remove that rhetorical escape. Once the data-generating mechanism is made transparent, the idea of model-free confirmation collapses.

2.4 Hosiasson-Lindenbaum and the Historical Perspective

Hosiasson-Lindenbaum proposed an early relevance condition: evidence confirms a hypothesis only if it is probabilistically relevant to it [Hosiasson-Lindenbaum, 1940]. This principle anticipated the core insight of the finite-case analysis: confirmation requires a substantive empirical link between evidence and hypothesis.

Her contribution was tragically cut short by her death in the Holocaust.[1] The subsequent literature shifted toward abstract priors and possible-world semantics, obscuring the relevance condition. The finite models restore its force.

3 Universal Quantifiers and Abstract Priors

The persistence of the paradox reflects a natural tendency to treat universal generalisations as ranging over unbounded domains. Abstract Bayesian priors—Dirichlet distributions, exchangeability assumptions—reinforce this tendency by analysing confirmation in terms of continuous proportions rather than discrete counts.

This perspective obscures distinctions that are sharp in finite settings. Independence blocks confirmation in every finite configuration; finite-resource depletion generates confirmation only under specific empirical structures. Continuous priors over unbounded spaces make these differences invisible, creating the impression that confirmation flows from logical equivalence itself. The finite models show that it does not. This resonates with the broader critique of model-free induction advanced by Norton [2003] and the emphasis on empirically grounded inquiry in Kelly [1996].

3.1 The Required Prior

The finite examples show that observing a non-raven confirms the universal only under a specific prior, namely one satisfying:

$$\mathbb{E}[P_N \mid P_R = 1] < \mathbb{E}[P_N \mid P_R < 1],$$

where P_R and P_N represent the proportions of black objects among ravens and non-ravens.

This requirement has been emphasised by Fitelson and Hawthorne [2010]. They show that non-black non-ravens confirm the universal only under priors that link the two colour-distributions, and defend these priors on grounds of Bayesian symmetry or exchangeability rather than empirical considerations. Their analysis reinforces the lesson of the finite models: the effect is not a logical consequence of equivalence but an artefact of a prior encoding the relevant empirical relationship.

[1] I am grateful to Richard Pettigrew for bringing her work to my attention.

4 Confirmation in Science: Natural and Social

The finite-case perspective aligns with a central lesson from natural science: confirmation is always model-relative. Newtonian gravitation is confirmed only because mass, distance, and acceleration are empirically linked; evolutionary biology identifies homology only because phylogenetic models specify how traits arise [Salmon, 1984, Woodward, 2003, Strevens, 2008].

Econometrics provides a parallel lesson. Sims's vector autoregressions were proposed as a "theory-free" alternative to structural macroeconometrics [Sims, 1980] and introduced in simple terms by Stock and Watson [2001]. But all interpretable VARs rely on strong structural assumptions. The Cholesky decomposition imposes a recursive causal ordering, and Kilian and Lütkepohl [2017] emphasise that VARs "are not theory-free": identification requires restrictions on timing, causality, or long-run structure.

5 Conclusion

The raven paradox persists only because confirmation has been analysed in abstraction from the empirical structure of finite domains. Once that structure is reinstated, the paradox disappears. Observing a non-raven confirms a universal generalisation only if the background assumptions encode a substantive relationship between categories. Without such assumptions, there is no confirmation.

More broadly, the discussion here illustrates a general lesson: confirmation is always model-relative. Attempts to reason in a model-free environment generate puzzles because, in the absence of structure, confirmation has nothing to operate on. The raven paradox is best understood as a demonstration of this principle.

References

Branden Fitelson and James Hawthorne. How bayesian confirmation theory handles the paradox of the ravens. In Ellery Eells and James H. Fetzer, editors, *The Place of Probability in Science*, pages 247–275. Springer, Dordrecht, 2010. doi: 10.1007/978-90-481-3615-5_11.

I. J. Good. The white shoe is a red herring. *British Journal for the Philosophy of Science*, 17(4):322–329, 1966. doi: 10.1093/bjps/17.4.322.

Carl G. Hempel. Studies in the logic of confirmation. *Mind*, 54(213):1–26, 1945. doi: 10.1093/mind/LIV.213.1.

Janina Hosiasson-Lindenbaum. On confirmation. *Journal of Symbolic Logic*, 5(4):133–148, 1940. doi: 10.2307/2267549.

Kevin T. Kelly. *The Logic of Reliable Inquiry.* Oxford University Press, Oxford, 1996. doi: 10.1093/0195078633.001.0001.

Lutz Kilian and Helmut Lütkepohl. *Structural Vector Autoregressive Analysis.* Cambridge University Press, Cambridge, 2017. doi: 10.1017/9781108164818.

Patrick Maher. Inductive logic and the ravens paradox. *Philosophy of Science*, 66(1):50–70, 1999. doi: 10.1086/392676.

Patrick Maher. Probability captures the logic of scientific confirmation. In Christopher Hitchcock, editor, *Contemporary Debates in the Philosophy of Science*, pages 69–93. Blackwell, Malden, MA, 2004.

John D. Norton. A material theory of induction. *Philosophy of Science*, 70(4):647–670, 2003. doi: 10.1086/378611.

Wesley C. Salmon. *Scientific Explanation and the Causal Structure of the World.* Princeton University Press, Princeton, 1984. doi: 10.1515/9781400854688.

Christopher A. Sims. Macroeconomics and reality. *Econometrica*, 48(1):1–48, 1980. doi: 10.2307/1912017.

James H. Stock and Mark W. Watson. Vector autoregressions. *Journal of Economic Perspectives*, 15(4):101–115, 2001. doi: 10.1257/jep.15.4.101.

Michael Strevens. *Depth: An Account of Scientific Explanation.* Harvard University Press, Cambridge, MA, 2008. doi: 10.4159/9780674034359.

Peter B. M. Vranas. Hempel's raven paradox: A lacuna in the standard bayesian solution. *British Journal for the Philosophy of Science*, 55(3):545–560, 2004. doi: 10.1093/bjps/55.3.545.

James Woodward. *Making Things Happen: A Theory of Causal Explanation.* Oxford University Press, Oxford, 2003. doi: 10.1093/0195155270.001.0001.