

Research Article

Mathew Coakley*

What Makes a Prediction Arbitrary? A Proposal

<https://doi.org/10.1515/opphil-2025-0067>

received January 17, 2025; accepted March 17, 2025

Abstract: What makes a prediction arbitrary? This article explores the possibility that one source of arbitrariness is asserting “P iff Q” when the justifications of P and Q are, in the relevant sense, independent. It uses this idea to draw a formal distinction between non-arbitrary and arbitrary predictions, even if they are also correct. It initially illustrates with Goodman’s New Riddle. There are, by now, so many different literature-prominent variants of grue, emerose or other odd-looking predicates that, surprisingly, the criterion handles them relatively straightforwardly. As the proposal is markedly distinct from existing resolution strategies centred on predicate projectability or counterfactual analysis, it sidesteps their well-known problems. The article makes the case that testing entailments for arbitrariness enjoys advantages over the currently dominant approach of directly putting conditions on when an F that is G should support “All F are G.” The article then surveys possible broader objections and offers two formal tools for identifying when and especially why a hypothesis makes good – that is, correct *and* non-arbitrary – predictions or not.

Keywords: induction, grue, non-arbitrariness, evidential inference, confirmation theory

1 Arbitrary Prediction and Justificatory Dependence

1.1 The Basic Thought

What makes a prediction *arbitrary*? It seems different to it being *incorrect*. Consider the following case. With a friend, you go to an area famous for having black kites and red kites. You agree to each write down a prediction about your day, of form “P iff Q.” You are sceptical that you will see a non-black kite, so you write “We will see a kite iff we see a black kite.”

You subsequently see several black kites and no other kites. There is a sense here in which your prediction was correct. You predicted “P iff Q,” and both P and Q occurred.

Your friend then shows you her prediction: “We will see a black kite iff Senegal win their World Cup qualifier.” You check the news, and indeed, they did. Under the previous logic, she also made a correct prediction. She can legitimately point out that her prediction could have been wrong – such as if Senegal had lost. But isn’t there still something wrong here? What does Senegal’s victory have to do with seeing a black kite (and *vice versa*)?

This article explores the following two thoughts. Both predictions were correct. Hers was also arbitrary. This might not be the only sense of the concept that is useful, but – according to this article – despite its relative simplicity, it is surprisingly useful.

The key idea is that even if a prediction of the form “P iff Q” correctly predicts both that P and that Q, if the reasons which currently justify believing P are independent of the reasons which currently justify believing Q,

* Corresponding author: Mathew Coakley, Independent Researcher, London, United Kingdom, e-mail: m.p.coakley.97@cantab.net

then to claim “P iff Q” is arbitrary. Roughly: the evidence (and background/*a priori* knowledge) justifying the black-kite-based proposition is independent of the evidence (and background/*a priori* knowledge) justifying the Senegal-victory proposition.

To forestall an immediate possible misunderstanding: arbitrariness relations hold between a particular prediction and the evidence. Thus, before any birds are seen, there is no necessary reason to think a black kite-based prediction is any more or less arbitrary than, say, a red kite-based one. It’s not about the predicates *per se*.

The proposal consists of three parts: (i) a criterion for when a theory makes a specific “P iff Q” style prediction, (ii) an account of correct prediction, and (iii) an account of arbitrary prediction. These are set out below, and this section concludes by noting the expository assumptions subsequently used to illustrate the approach.

1.1.1 Making a Biconditional Prediction

Throughout this article, any hypothesis H will be said to “make” the prediction “P iff Q” if:

- (a) H entails “P iff Q”; and
- (b) H does not entail “P & Q” and does not entail “not-P & not-Q.”

I’ve come to believe that (b) can be acceptably, and perhaps desirably, dropped. However, if so, some results are somewhat counterintuitive and orthogonal to the key insights. Thus, this article adopts both (a) and (b).

1.1.2 Making a Correct Prediction

Let “P iff Q” be a **correct** prediction iff either “P and Q” is true or “not-P and not-Q” is true.

The idea underlying the notion of “correct prediction” is neither original nor solves any of the problem cases later discussed. It is used to distinguish correctness from non-arbitrariness and thus identify when these coincide and when they do not.

Trivially, for completeness, also let “P iff Q” be an incorrect prediction if either “P and not-Q” is true or if “Q and not-P” is true. This specific formulation does no significant work in what follows and can be alternatively captured in a range of commonplace ways.

1.1.3 Making an Arbitrary Prediction

The previous example of arbitrariness used a very general reason-based notion of justificatory dependence; thus, a more specific working definition is required to further concretely explore the underlying idea. Here is the one this article will use:

Justificatory dependence: the justification of P depends¹ upon the justification of Q - hereafter j(P)/j(Q) - iff every set of those justifiably believed things currently sufficient to justify that P also justifies that Q.

In other words, you couldn’t justifiably doubt Q, based on doubt over elements of what you currently believe, without simultaneously doubting P.

If the justification of S is not dependent on the justification of R, and *vice versa*, and both S and R are justified, then the justification of S is *independent* of the justification of R.

¹ The use of “depend” and “dependent” is not an appeal to any sort of Cambridge dependence, nor to a dependence logic more generally. They are used throughout this article only and always as defined.

This allows us to spell out the previous distinction between correct prediction (roughly: getting a prediction right) and arbitrary prediction (roughly: getting it right, but not for the right reasons).

Let “ P iff Q ” be an **arbitrary** prediction if the justification of P is independent of the justification of Q .²

Arbitrariness also holds across logical equivalences: if “ S iff T ” is arbitrary and both entails and is entailed by “ P iff Q ,” then “ P iff Q ” is arbitrary. Given the definition of justificatory dependence, if P entails Q , and the belief that P is justified, then necessarily $j(P)/j(Q)$. Thus, for any two justified propositions P and Q , if neither proposition entails the other, then *one and strictly only one* of the following three things must always obtain: (i) $j(P)/j(Q)$, (ii) $j(Q)/j(P)$, or (iii) neither $j(P)/j(Q)$ nor $j(Q)/j(P)$. “ P iff Q ” is arbitrary if (iii).

1.1.4 Expository Assumptions

More broadly, to talk of “justifying things” already imports much disagreement over how propositions are justified. Fortunately, and somewhat surprisingly, this issue can be mainly put to one side in what follows. It is orthogonal – though very important – as the account here supervenes upon, but does not require, any specific theory of proposition justification, although it will need to assume that propositions, when justified, provide epistemic access to those justifications.

Although a specific theory of proposition learning is not required in principle, some rough or implicit account is needed in practice to illustrate the approach concretely. Thus, to set out the argument, this article will assume two things. First, it is possible to justifiably believe in the existence of a set of evidence. Second, it is possible for a set of evidence to “establish” some propositions about the universe; specifically, that justified belief in the existence of that set of evidence, potentially combined with some *a priori* and/or background knowledge, justifies belief in, or acceptance of, the propositions.

These assumptions can be relaxed, and the overall argument stands. Furthermore, exactly how evidence establishes a proposition (such as “there exists an F ”) will not matter, with the term “establish” intended to permit rival accounts. Of course, it will matter that a compelling theory permits this, and debates over the best account as such are very important but orthogonal to this article’s core arguments.³

Given these assumptions, in empirical contexts, therefore:

For any set of evidence that establishes that P and that Q , the justification of P depends on that of Q – $j(P)/j(Q)$ – iff every subset of the evidence that establishes that P also establishes that Q .

Section 1.1 in summation:

- Hypothesis H makes a **correct** prediction if H entails “ P iff Q ,” does not entail “ $P \ \& \ Q$,” does not entail “not- $P \ \& \$ not- Q ,” and it is true that “ $P \ \& \ Q$.”
- This prediction is **arbitrary** if neither $j(P)/j(Q)$ nor $j(Q)/j(P)$, namely, there is a subset of the evidence that establishes P that does not establish Q , and *vice versa*.

1.2 Hypotheses Making Non-Arbitrary and Arbitrary Predictions

You hear your village clock strike 10 am, see a swan a few minutes afterwards, and then, walking closer, see it is white. This evidence establishes, among other things, the also correct proposition that “There exists a white swan in the morning.” Here are two hypotheses:

² This condition and the underlying broader logic were first proposed to cover beliefs in general – normative and empirical – see the discussion in Coakley, *Motivation Ethics*, 215–33. However, this article applies the ideas solely to evidential contexts.

³ So, to take simple illustrative examples, the evidence for a sense-datum theory would consist of a set of sense data. In contrast, for an early adverbial theory, it would be a set of adverbial perceptual engagements. Chisholm, *Perceiving*; Aune, *Knowledge, Mind and Nature*; Sellars, *Science and Metaphysics*.

H1: (All swans in the morning are white) & (all swans not in the morning are not-white).
 H2: (All swans are white).

Both make multiple correct predictions, such as:

(there exists a swan in the morning) iff (there exists a white swan in the morning).

However, H1 makes an arbitrary prediction:

(there exists a white swan) iff (there exists a swan in the morning).

There is evidence establishing that there exists a white swan that doesn't establish that it is a swan in the morning, and *vice versa*. For instance, your visual evidence about the bird's shape and colour, along with background knowledge of bird species, is sufficient to establish it's a white swan, but that doesn't also necessarily establish it's in the morning, even though you do justifiably believe it's the morning. Hence:

- (i) It is not that $j(\text{there exists a white swan})/j(\text{there exists a swan in the morning})$; and
- (ii) It is not that $j(\text{there exists a swan in the morning})/j(\text{there exists a white swan})$.

The issue is not that H1 correctly predicts that "P iff Q" where neither P nor Q entails the other. H2 has many such predictions, and they are non-arbitrary, such as:

(there exists a white swan in the morning or there exists a non-white swan) iff (there exists a swan in the morning).

Every subset of the evidence that establishes that "There exists a white swan in the morning or there exists a non-white swan" also establishes that "There exists a swan in the morning" hence $j(\text{There exists a white swan in the morning or there exists a non-white swan})/j(\text{there exists a swan in the morning})$.

More generally, "All F are G" will always make some correct predictions given evidence establishing that there exists an F that is G. Whether these are all non-arbitrary will depend on the particular F and G in question. The rest of this article explores the implications. Section 2 applies the approach to grue-like cases. Section 3 contrasts it with rivals. Section 4 considers structural worries and resources. Section 5 considers case-based objections, such as where an apparently good hypothesis is judged arbitrary, or a seemingly bad one judged non-arbitrary.

2 Grue-Like Cases and the Failings of Induction-as-Uniformity

Goodman's challenge to a formal account of induction was to distinguish when a green emerald observed before T should confirm "All emeralds are green," but not "All emeralds are grue," if something like the following mutual entailment symmetry holds: "There exists a grue emerald observed before T" $\leftrightarrow$ "There exists a green emerald observed before T."

There are multiple suggested formulations of the predicate grue. Whichever is used, "All emeralds are grue" makes an arbitrary prediction; "All emeralds are green" does not.

2.1 Time-Indexed Grue Formulations⁴

Assume we observe a green emerald before some future time T, this establishing that "there exists a green emerald before T."

⁴ Goodman's initial pre-Grue marble-based description of the problem matches this type of formulation. Goodman, "A Query on Confirmation," 383. This was also common in early discussions, for instance Achinstein and Barker, "On the New Riddle of Induction."

x is grue1 iff (x is green and before T) or (x is blue and not-before T)
 x is green1 iff (x is green and before T) or (x is green and not-before T)

“All emeralds are grue1 ” makes an arbitrary prediction, namely that:

(there exists a green1 emerald) iff (there exists an emerald before T).

There is a subset of the evidence sufficient to establish that “There exists a green1 emerald” that doesn’t also establish that “There exists an emerald before T ” – e.g. our observation of its greenish hue and evidence of those features that distinguish an emerald from other possible options. There is a subset of the evidence sufficient to establish that “There exists an emerald before T ” that doesn’t also establish that “There exists a green1 emerald” – e.g. our observation of our watch and emerald shape etc. Hence, the prediction is arbitrary, in the same way that “We will see a black kite iff Senegal win their World Cup qualifier” was arbitrary. The two propositions on either side of the iff clause have separate evidence establishing them.

We could, for instance, be sceptical about the existence of the evidence establishing that it is before T – maybe we have misread the calendar, or our watches are wrong – yet not be sceptical that the emerald is green1 . Whereas, in the superficially symmetric case, “All emeralds are green1 ” makes a non-arbitrary prediction:

(there exists a grue1 emerald) iff (there exists an emerald before T)

Any set of evidence sufficient to establish that there is a grue1 emerald also establishes there is an emerald before T . If we were sceptical about the evidence that it was before T – maybe we have misread the calendar, or our watches are wrong – we should be sceptical as to whether it is a grue1 emerald.

The critical point is not that “All emeralds are grue1 ” gets a prediction wrong. Along with “All emeralds are green1 ,” it correctly forecasts that the emerald observed before T will be green. Instead, it makes an arbitrary prediction about what it gets right, namely, that our warrant for believing it green requires the evidence that it is before T . Goodman’s genius was to spot that correct prediction isn’t enough. This article agrees. Prediction needs to be correct *and* non-arbitrary.

2.2 Observation-Indexed Grue Formulations (Jackson’s “Problem Case”)

Jackson famously argued⁵ that the really problematic of the grue definitions was instead the following:

x is grue2 at t iff (x is green at t and observed before T) or (x is blue at t and not observed before T)
 x is green2 at t iff (x is green at t and observed before T) or (x is green at t and not observed before T)

The substantive result still holds. “All emeralds are grue2 ” makes an arbitrary prediction:

(there exists a green2 emerald) iff (there exists an emerald observed before T).

There is a set of evidence sufficient to establish that “there exists a green2 emerald” that doesn’t also establish “there exists an emerald observed before T ” and *vice versa*.

With the greenish parallel case, however, “All emeralds are green2 ” makes the prediction that:

(there exists a grue2 emerald) iff (there exists an emerald observed before T).

Here, every set of evidence that establishes “there exists a grue2 emerald” also establishes that “there exists an emerald observed before T .”

⁵ Jackson, “Grue.”

2.3 A Generic Grue-Style Case

Grue cases come in two distinct types:

- o For “exclusionary” grue predicates, an object is grue iff it is [(green and F) or (blue and not-F)].
- o For “inclusionary” grue predicates, an object is grue iff it is [(green and F) or (blue and G)] where the object, in principle, could be *both* F and G.

An example of an exclusionary variant would be grue1 or grue2: an object is either “observed before T” or “not observed before T”; it cannot be both. An inclusionary example would be GrueMon, where something is GrueMon iff it is [(green and exists on a Monday) or (blue and exists on a non-Monday day)]. A green emerald existing on a Monday could also exist on a Thursday.

Exclusionary and inclusionary grue predicates have different underlying logical structures. Some prominent proposals only address one (Section 3.3).

Generic exclusionary cases.

Let the predicate “OT” generically represent “in our sample,” “observed before T,” “observed at t before T” or any similar predicate whereby:

x is grue* iff (x is green and OT) or (x is blue and not-OT).
 x is green* iff (x is green and OT) or (x is green and not-OT)⁶.

“All emeralds are grue*” makes an arbitrary prediction about a green OT emerald:

(there exists a green* emerald) iff (there exists an emerald that is OT)

By contrast, “All emeralds are green*” will not:

(there exists a grue* emerald) iff (there exists an emerald that is OT).

Generic inclusionary cases

Let M and M* be distinct predicates such that an object could, in principle, be both but need not, and where evidence of both is independent of colour judgements. For example M is “exists on a Monday,” and M* is “exists on a non-Monday.”

All emeralds are [(green and M) or (blue and M*)] makes an arbitrary inference about a green M emerald, namely:

(there exists a green emerald) iff (there exists a not-M* emerald).

The evidence that it is green is independent of the evidence that it is not M*, e.g. that it exists on a Monday.

2.4 Using Grue/Bleen as the Primitive Terms

Nothing in the results hangs upon which terms – grue/bleen or green/blue – we use to define the hypotheses or evidence. To illustrate, here is a generic case where (i) x is grue iff (x is green and OT) or (x is blue and not-OT); and (ii) x is bleen iff (x is blue and OT) or (x is green and not-OT). The relevant predicates in grue/bleen terms are thus:

x is grue* iff (x is grue and OT) or (x is grue and not-OT).
 x is green* iff (x is grue and OT) or (x is bleen and not-OT).

⁶ The generic formulation seems to capture Goodman’s favoured usages, for example notably Goodman, *Problems and Projects*, 359; or Goodman, *Fact, Fiction and Forecast*, 74.

“All emeralds are grue*” still makes an arbitrary prediction, namely:

[there exists an emerald that is ((grue and OT) or (bleen and not-OT))] iff [there exists an emerald that is OT]

There is evidence establishing that the emerald is “(grue and OT) or (bleen and not-OT)” that doesn’t establish that the emerald is OT: the evidence of its visual appearance is sufficient. There is evidence establishing that the emerald is OT – say our observation of our watch, or clock, or that it’s observed before some particular time etc., whatever the OT is – that doesn’t establish it’s also ((grue and OT) or (bleen and not-OT)).

By contrast, “all emeralds are green*” makes a symmetric but non-arbitrary prediction, *even if expressed in grue/bleen terms*, namely:

(there exists an emerald that is ((grue and OT) or (grue and not-OT))) iff (there exists an emerald that is OT).

Here, any set of the evidence establishing that the emerald is ((grue and OT) or (grue and not-OT)) also establishes that the emerald is OT.

These results are unsurprising: arbitrariness relations hold across logical equivalences. It is irrelevant which lexical terms describe the predicates and hypotheses.

3 Rival Extant Approaches and How They Differ From Here

The grue literature is truly vast. For example, more than a quarter of a century ago, Stalker edited a book-length collection of diverse proposals,⁷ and since then, the papers and proposals have proliferated.

There is, however, a general dominant resolution approach. And it is a rival to that of this article. Namely:

The predicate-condition strategy: a condition is provided for when “All F are G” is supported by an F that is G, or an observation or belief thereof, and also for when it is not.

Theories adopting this strategy come in one of three types:

- (i) *Qualitative-predicate theories*, where grue is said to be different in kind to green.
- (ii) *Counterfactual theories*, which appeal to a relevant counterfactual to break the apparent grue/green symmetry.
- (iii) *Agent-relative theories*, where it is either how or why a particular agent forms conclusions about the green/grue emerald that matters for predicate projection.

Each has distinct strengths and challenges, as noted below. They also all potentially have some common problems and weaknesses, as discussed subsequently. This article offers no such predicate-generalisation condition. Instead, it is, in effect, a hypothetico-deductive proposal that evaluates entailments by an arbitrariness test. Both strategies can be applied to the same cases but differ radically in how they approach the task. Key differences are noted below.

3.1 Qualitative-Predicate Theories

Simply on the face of it, grue seems suspect, made-up or, in a sense, not real. Qualitative-predicate theories aim to explain why. Prominent suggestions are that green – but not grue – is a natural kind,⁸ better entrenched,⁹ or genuine, not “pseudo.”¹⁰ Of course, the famous problem here is that the grue-sympathiser will make the opposite judgement; thus, the challenge is to show why their belief is unfounded.

⁷ Stalker, *Grue!*.

⁸ Quine, “Natural Kinds.”

⁹ Goodman, *Fact, Fiction and Forecast*.

¹⁰ Shoemaker, “Properties, Causation and Projectibility;” Bealer, *Quality and Concept*; Armstrong, *What is a Law of Nature?*

In terms of strategy, this article's entailment-arbitrariness approach is different in type: there are no assumptions about predicates having any relevant qualitative features. To the extent that qualitative predicate theories assume what needs to be shown, that is a relative advantage.

3.2 Counterfactual Theories

Under a prominent *grue* case, an object at a given time is *grue* if examined and green or unexamined and blue. But most of us don't – usually – think examining objects changes their underlying properties. Hence, Jackson's initial famous counterfactual resolution¹¹ relied, in a good inductive generalisation to “All *F* are *G*,” on us being entitled to assume that, if an *F* that is *H* and *G* had not been *H*, it would still have been *G*.

Again, this is a very different type of resolution from here, where there are simply relations between the justifications of propositions. The entailment-arbitrariness approach has no counterfactual and does not implicitly appeal to any assumptions about nearest worlds or similar. The facts it draws upon are facts about this world: the structure of the set of evidence an agent has.

Thus, it sidesteps the well-known core complaint that counterfactual resolutions smuggle in the conclusion they aim to establish.¹²

3.2.1 The Reoccurring Ghost of the Problem Case

In addition to general worries about the strategy, Jackson's proposal encountered a response that would haunt most such conditions (both counterfactual and agent-relative): that of the problematic unaddressed case.

Jackson's initial counter-factual suggestion, for example, is vulnerable to a green OT emerald confirming “All *emerose1* are green” (where again OT covers the relevant qualifier, such as observed, examined, observed at *t* before *T*, etc.):

x is *emerose1* iff (*x* is an emerald and OT) or (*x* is a rose and not-OT).

If a green *emerose1* hadn't been observed, it would still have been green. So, under the initial proposal, the generalisation is acceptable. Partly as a result, Jackson ended up changing his counterfactual condition – with Pargetter¹³ – to one that blocks this sort of confirmation. In response, a further problem case was raised (see subsequently).

This raises a central methodological issue that will resurface in this article: Are people's intuitions about the cases actually reliable?

3.2.2 Are Emerose Cases Always Wrong to Project?

The extant literature generally takes it as self-evident that “All (observed emeralds or unobserved roses) are green” being supported by an observed green emerald is a “bad induction” as Godfrey-Smith¹⁴ terms it, or as Jackson and Pargetter explain “we do not normally expect unexamined *emeroeses*, that is, unexamined roses, to be green on this account.”¹⁵

¹¹ Jackson, “Grue.”

¹² For example, Roskies, “Robustness and the New Riddle Revived.”

¹³ Jackson and Pargetter, “Confirmation and the Nomological.”

¹⁴ Godfrey-Smith, “Goodman's Problem and Scientific Methodology,” 576.

¹⁵ Jackson and Pargetter, “Confirmation and the Nomological,” 426.

Everyone seems to endorse these judgements, and subsequent proposals were thus designed and crafted to block the confirmation. But is that correct? It seems potentially not. To see why, here is the hypothesis phrased as a logical equivalence:

“All emero₁ are green” $\leftrightarrow$ “(All OT emeralds are green) & (All not-OT roses are green)”.

A green OT emerald offers intuitive support for the first – “All OT emeralds are green” – clause. Whether the second clause is supported should depend upon whether we have any evidence of the colour of roses. Imagine we do not. Well, then, we’ve observed an object being green, which seems relevant evidence that other objects will be green. *Purely on its own*, a green emerald should support the hypothesis that all unobserved roses are green. The “All emero₁ are green” hypothesis is problematic, but for a different reason.

What seems bad, presumably, about “All emero₁ are green” being confirmed is the background evidence that there are (OT) roses that are non-green. However, if so, “All emero₁ are green” makes an arbitrary prediction about *that* evidence:

(there exists a non-green rose) iff (there exists an OT rose)

Instead, imagine that we have no evidence of the existence of a non-green rose – either all OT roses have been green, or we have no evidence about rose colours. If so, “All emero₁ are green” being confirmed by a green OT emerald would not seem obviously problematic.

If, in contrast, we do have such evidence, “All emero₁ are green” straightforwardly makes an arbitrary prediction. It is the non-green OT rose, not the green OT emerald, which should count against “All emero₁ are green.” Getting this wrong will lead to conditions which either rule out too much or fail to rule out the right hypothesis–object pair.

Now consider, via Godfrey-Smith, emero₂:

x is emero₂ iff (x is an emerald) or (x is a rose and not-OT).

This was a problem for Jackson and Pargetter’s revised proposal. But, with the same logic above, if we have no evidence of non-green roses, then a green OT emerald confirming “All emero₂ are green” is not problematic. The hypothesis will, however, make an arbitrary prediction about any non-green OT rose, namely:

(there exists a non-green rose) iff (there exists an OT rose).

This article’s test for non-arbitrariness produces results – notably with emero₂ cases – conflicting with received wisdom from the literature. The latter may just be mistaken. In both emero₁ and emero₂ cases, it is evidence of the non-green colour of roses – *not* the green emerald – that should count against the generalisation.

The main structural challenge for counterfactual approaches is establishing the counterfactual without presupposing that *grue* is a bad predicate (which would be a qualitative-predicate approach and potentially beg the question). This article’s approach does not face that challenge simply by being based on entailments and the structure of the evidential set.

In addition, the emero₂ cases illustrate the danger of relying on intuitions when terms seem weird, but the overall induction might be justified. Because it is general and has a clear logic, the entailment arbitrariness approach can explain why cases are good or bad inductions. This is independent of our localised reactions to odd-looking hypotheses or predicates. That is an advantage.

3.3 Agent-Relative Approaches

For an agent-relative theory, the projectability of a predicate depends upon some aspect of how or why the agent learns or knows that the object is that predicate.

Almost all of the extant suggestions offer a specific condition to rule out bad inductions (with a notable exception being Warren, “*The Independence Solution to Grue*,” where there are multiple distinctly motivated conditions).

Before surveying the strengths and weaknesses of these sorts of approaches, it’s worth noting that Warren offers a careful theory history of the overall strategy, tracing it back initially to Wilkerson,¹⁶ Moreland¹⁷ and Jackson,¹⁸ with more recent examples being Okasha,¹⁹ Godfrey-Smith,²⁰ Schramm,²¹ Freitag,²² and of course now Warren.²³

Warren calls the overall approach the “independence strategy.” All these authors are trying to draw a distinction where the grue predicate *depends* on something that the green predicate does not, so ruling out generalising features that depend on that thing rules out “all emeralds are grue” but not “all emeralds are green.” Hence, green can be projected because it is “independent” in the relevant sense. Everyone disagrees about the content of that sense (but typically uses terms like independence or dependence to describe its absence or presence).

My terminology here separates counterfactual from agent-relative approaches because, for proponents, the counterfactual is meant to be a feature of the world (or of the nearest possible world), one that is common to all agents. Whether this is correct is one challenge. Sceptics worry it is actually agent-relative, and grue-sympathisers thus don’t have reason to accept the green-preserving counterfactual.

As Warren usefully highlights, however, there is an underlying common approach. This article’s proposal is very different. To illustrate, take the two most recent cited proposals.

Freitag in “*I’ll Bet You Solve Goodman’s Riddle*” introduces the core idea as a dialogue, and subsequently in “*The Disjunctive Riddle and the Grue-Paradox*” sets it out descriptively. There are two key theory moves. First, a predicate P is “epistemically discriminating” with regard to two sets if an agent knows that (i) all the elements of the first set are P, and (ii) all of the second set are not-P.²⁴

In other words, “Discriminating predicates are such that it is known by S that their extension covers all the samples and only them.”²⁵ “Being in our sample” would, as such, necessarily be discriminating: the feature cannot be projected to the set of objects outside it; it is directly defeated.

Second:

An (inductively confirmed) hypothesis is derivatively defeated if and only if the pertinent inductive evidence epistemically depends on the inductive evidence for the projection of a discriminating predicate.²⁶

So, for instance, imagine that all sampled emeralds have been green. If grue_S is “(green and sampled) or (blue and not-sampled),” then they have also all been grue_S. Thus, “all emeralds are green” and “all emeralds are grue_S” both describe all sampled set members. This is the problem.

However, “being sampled” is an epistemically discriminating predicate: all the sampled emeralds are it, and all the unsampled emeralds are not. The agent believes that the emeralds are green-coloured not because they are part of the sample but because of visual considerations. Green can be projected. However, they must know the emeralds are in the sample to know they are grue_S. So, the projection of grue_S is derivatively defeated.

¹⁶ Wilkerson, “A Gruesome Note.”

¹⁷ Moreland, “On Projecting Grue.”

¹⁸ Jackson, “Grue.”

¹⁹ Okasha, “What Does Goodman’s ‘Grue’ Problem Really Show?”

²⁰ Godfrey-Smith, “Goodman’s Problem and Scientific Methodology;” and Godfrey-Smith, “Induction, Samples, and Kinds.”

²¹ Schramm, “Evidence, Hypothesis, Grue.”

²² Freitag, “I’ll Bet You Solve Goodman’s Riddle;” and Freitag, “The Disjunctive Riddle and the Grue-Paradox.”

²³ Warren, “The Independence Solution to Grue.” Another example is Coakley, “Hempel, Grue and the Logical Empiricist Baseline.”

²⁴ Freitag, “The Disjunctive Riddle and the Grue-Paradox,” 187.

²⁵ Ibid., 188.

²⁶ Ibid., 194.

One attractive feature here is that Freitag offers an insightful justification for *why* the grue generalisation should be ruled out. If knowing something was F is necessary to conclude it is P, then it seems illegitimate to project the P-ness to known non-Fs.

One complaint – via Dorst²⁷ – is that the resolution begs the question by relying on assumptions, notably over language or counterfactuals, that a grue-sympathiser would not accept. Freitag in reply²⁸ makes the case that the content of the agents' beliefs, not how they are labelled, ultimately matters.

Freitag only discusses a particular type of grue case, which mirrors those of grue1 and the generic grue* from earlier. His resolution and that of this article here potentially agree, and *for the same reasons*: why the agent knows the object is grue or green is key.

In other words, all his discussed cases are where grue is [(green and O) or (blue and not-O)], and the relevant evidence is all green and O. The predicate O here is, therefore, discriminating. This is also the case that captures the motivating intuition: if knowing the object is O is necessary to know it is grue, it does seem illegitimate to project its grueness to known not-Os.

Like almost every proposed condition in the literature, the article cites a case where it gets the answer right and for perhaps the right reasons. As is common with those other suggestions, the problems arise from the unaddressed cases.

First, Freitag's proposal only appears to cover "exclusionary" grue predicates. Consider again "All emeralds are GrueMon" where an object is GrueMon if it is [(green and exists on a Monday) or (blue and exists on a non-Monday day)]."

Existing on a Monday (rather than a non-Monday) is not an epistemically discriminating predicate. We do not know that objects that exist on a Monday do not also exist on a Thursday. So, the generalisation appears not defeated by a green emerald that exists on the Monday. This problem seems to apply to *any* grue that is [(green and A) or (blue and B)] where the object could potentially be A and B. We might even know it is B – e.g. it does exist on a Thursday – so that set is not empty.

Second, it appears the proposal implicitly endorses *all* possible emerald-style predicates. "All (emeralds or unobserved roses) are green" could, as such, be generalised from a set of green emeralds to a set of unobserved roses. None of the Freitag papers explicitly address the emerald cases, so it is hard to be sure if the best version of the proposal would adequately cover them. The condition seems to establish that a set of observed x that are all green can have that *greenness* projected to a set of unobserved x (but not their grueness). If so, a green emerald would support "All (emeralds or unobserved-Y's) are green" for every possible Y.

Third, as a result, it seems vulnerable to simply reintroducing all the problems via the object. Consider:

An object is "operald" iff it is (an opal or an unsampled emerald).

A set of sampled blue opals would seemingly support the hypothesis that "all operalds are blue" and thus that all unsampled emeralds are blue – substantively precisely the case previously said to have been defeated. Sampled green emeralds do not defeat the hypothesis, so if these are included, "all operalds are blue" remains an apparent good induction.

Because only one type of grue-like case – All emeralds are (green and F) or (blue and not-F) – is fully set out, it's hard to know if some amendment, generous interpretation or workaround can address the above problem cases.

Thus, although Freitag's discriminating predicate/dependence theory and this article agree on the pure exclusionary grue cases, they potentially diverge quite radically in other cases, which is unsurprising given their very different approaches.

The underlying problem with the discriminating-predicate approach is that there are many ways of introducing arbitrary qualifications to hypotheses, and using the twin notions of discriminating predicates

²⁷ Dorst, "Bet Accepted."

²⁸ Freitag, "Why Doxastic Dependence Defeats Grue."

and epistemic dependence only blocks some of them. This weakness stems from the approach's virtue: it is relatively parsimonious.

3.4 There is a Pattern Here...

Both counterfactual and agent-relative approaches standardly propose a condition C. The argument is that “All emeralds are grue” fails it, and “All emeralds are green” does not. This is illustrated via a particular type of grue case.

Someone then points to a different case, which C seemingly gets wrong. The condition is then reformed, and a new problem is raised, or the debate moves on. Anyone with even a cursory knowledge of the grue literature will recognise this repeating pattern. What's interesting here, however, is that even if not gaining traction, the conditions often contain real insight.

Jackson, for example, spots something correct and important. Usually, we think an object observed to have an underlying property would still have had that property even if, counterfactually, we hadn't observed it.

Similarly, Freitag persuasively argues that if knowing an object was P was necessary to know it was F, then projecting that to similar objects known to be non-P is illegitimate.

And yet, these sorts of promising and insightful proposals subsequently reliably run into challenge cases. This is not a coincidence. The grue literature – diverse, with multiple suggested resolution strategies, and with proposals historically vulnerable to ingenious new problem predicates – looks precisely as would be expected if the underlying difficulty was the introduction of arbitrariness *and where this could be done in many different places and ways*.

Counterfactual approaches are vulnerable to arbitrary predicates where the counterfactual is not relevant. Agent-relative approaches are vulnerable to cases where the particular form of agent-relativity is not the source of the problems.

The rival approach of this article – using entailments to test for arbitrariness – effortlessly circumvents this problem because every time an arbitrary qualification is added anywhere to the theory, it always shows up in its entailments.

3.5 Why Not Have Multi-Conditions?

Since Goodman's initial discussion, many specific predicate conditions have been proposed, and reliably – even if they get an important case correct – they then run into problems with other types of cases. Warren²⁹ aims to correct for this sort of failure.

Warren takes previous examples of the independence strategy and considers which cases they get right and wrong, and therefore proposes *multiple* distinctly motivated conditions that, in combination, aim to block the problem generalisations.

Consider that some F's are sampled to see if they are G, with the question being whether “All F are G” is a warranted conclusion in the relevant sense. Having noted that the sample needs to be large enough, the article proposes three additional independent conditions:

First, the methods of sampling “MS” and of observation “MO” need to not *themselves* affect the results obtained.

²⁹ Warren, “The Independence Solution to Grue.”

Methods *MS* and *MO* are methodologically independent of *F* and *G* if and only if neither the *F*-population nor the distribution of *G* over the *F*-population would be altered by applications or non-applications of *MS* or *MO*, in relevant situations considered as actual.³⁰

Second, what is sampled must not be unduly biased compared to what is not.

Method *MS* is partition independent over *F* if and only if for any *F*, there is a minimal *MS*-variant world where that *F* is sampled by *MS*.³¹

Finally, sometimes other evidence or things believed about the evidence should block or undermine an induction (an insight from Freitag). So in an epistemic context *c* and making an inductive inference *I*:

I is well-based in *c*, if and only if *I*'s premises do not inferentially depend on (the premises of) an inductive argument, *I**, that is blocked or otherwise undermined in *c*—it is non-well-based otherwise.³²

What appears attractive here is that, by offering up multiple conditions, a greater range of seemingly problematic confirmations are blocked. Moreover, each of the proposed conditions does seem to capture something important.

The chief difficulty is: what if our intuitions about the cases are wrong? Then, we will either adopt a bad principle – being too narrow or too broad – or adopt the wrong explanation.

Consider, for instance, Warren's ingenious solution to one type of emeroses case:

Recall that something is an emeroses2 just in case it is either an emerald or an unobserved rose. I don't think this slips through the methodological independence net. Even if it does though, that is fine. We can say that it is of a kind with a third variant—something is an emeroses3 just in case it is either an emerald or a rose. This definition is not relativized to sampling or observations at all. Yet if we observe only emeralds, find them all green, then try to project "all emeroses are green", we run into trouble in the exact same way we do with "emeroses2". The difference is that with "emeroses3" our having an unrepresentative sample could have been avoided by observing some roses. Not so with "emeroses2".³³

One obvious worry is: what is it to be "of a kind"? That seems to be doing all the work, and the sceptic would protest that unobserved roses are not "of a kind" with roses in general. Second, what if roses are extremely rare? Our sample will not now necessarily be unrepresentative of the population of emeroses3.

But imagine those worries are overcome. There's a more fundamental problem. Why try to bend conditions to match our case-based intuitions if they themselves are unreliable? Recall: if we haven't observed non-green roses, "all emeroses2 are green" *should* be supported by a green emerald.

After all, it's the conjunction of two hypotheses: "All emeralds are green" and "All unobserved roses are green." The first is straightforwardly supported. But so too – absent other evidence – is the second: we've observed a green object, which makes us more confident other objects are green. That's a desirable result. Conditions with "All emeroses2 are green" supported by a green emerald on its own shouldn't be actively tweaked or generously interpreted to block this.

It's an apparent strength of Warren's approach of combining multiple independent criteria that it is flexible enough to be adjusted or amended to meet new or unanticipated problem cases. Indeed, this is an explicit goal. "Some accounts rule out too much, others too little. Like Goldilocks, we need an account that is instead *just right*."³⁴

Thus, the overall worry is not interpretative, such as what the three conditions really entail. It concerns the entire problem-case whack-a-mole strategy. Not all weird-looking cases are necessarily problematic. Over-deference to intuitions or conventional wisdom risks modifying or interpreting conditions so that the desired results are delivered and the wrong conditions are adopted.

³⁰ Ibid., 1316.

³¹ Ibid., 1318.

³² Ibid., 1322.

³³ Ibid., 1317–8.

³⁴ Ibid., 1315.

Furthermore, some cases are also just intuitively unclear, especially if there is no past literature judgment to which to defer. For instance, let something be “rainbow pink” iff it is generally pink but also has a rainbow stripe across it. In this world, should we believe that “All unicorns are rainbow-pink” because we keep observing green emeralds, grey shoes, brown sticks, red foxes...? The – probably counter-intuitive – correct answer, by the way, is yes. But the interesting question is: why? Systemising intuitions won’t provide such an explanation.

With case-focussed predicate-condition theories, it matters greatly whether our judgements about the cases are sound. In contrast, the entailment-arbitrariness approach of this article doesn’t ultimately defer to our intuitions. It’s informative when it conflicts with them (Section 5). But the case for it is much more general.

To believe “P iff Q,” it’s not enough that P and Q. The biconditional itself needs justifying. It’s a different sort of approach.

In content, it is also much more parsimonious than a triple-condition, where each has multiple moving interpretable theory parts. Under the entailment arbitrariness approach, there is a simple test:

Does the hypothesis entail a “P iff Q,” but not (P and Q) and not (not-P and not-Q), where the justification of P and Q are independent?

If so, it makes an arbitrary prediction. If not, not. That’s it.

3.6 A Generic Problem for Predicate-Condition Approaches

Predicate-condition approaches have “All F are G” supported by an F that is G (or observations/evidence thereof) iff condition C is also met. They disagree as to C.

Because the classic problem cases – *grue/blite/emeroose* – involve an F that is G where most believe that “All F are G” should *not* be supported, C typically puts further constraints on generalisations, so only a subset of possible F’s that are G are generalised. Consider, though, a classic *emeroose*-style hypothesis:

H1: “All (roses or not-OT emeralds) are red”.

Now, take a green OT emerald. The object is not a [(rose or not-OT emerald) that is (red)]. For predicate-condition approaches, it doesn’t seem they have to make a problematic confirmation here.

But there is a problematic omission. H1 should be *undermined* by the object/observation. Blocking bad confirmations is not enough. Some theories should be undermined by an object that does not logically refute them and is not an instance of their generalisation.

Nor is simply appealing to the absence of positive confirmation persuasive. H1 does worse here than H2 – “All roses taste of chicken” – and a green OT emerald that has not been tasted. H2 seems genuinely neutral. H1 implies not-OT emeralds will be red, and yet all emeralds so far have been green. It does worse.

The entailment-arbitrariness approach gets this case correct without having to gerrymander on extra conditions. H1 entails (there exists a non-red emerald) iff (there exists an OT emerald). This is arbitrary.

Moreover, beyond intuitions, the approach explains what’s wrong with the hypothesis: it arbitrarily ties non-red emerald status to time status.

4 Challenges and Advantages Specific to the Entailment-Arbitrariness Approach

The entailment-arbitrariness approach has no special role for predicates and does not view the task of a theory of induction as putting conditions on generalisations. It covers predicates and generalisations, but focuses on entailments of the form “ P iff Q ” and the justifications of P and Q . To wit: the hypothesis “There exists a fox iff there exists a pink badger” is directly scrutinised, not clumsily shoe-horned into an “All F are G ” framework. This creates relative strengths and challenges.

4.1 Shouldn't Bi-Conditionals Have Bi-Conditional Justificatory Dependence?

Shouldn't biconditionals of form “ P iff Q ” be justified by mutual justificatory dependence, namely that “ P iff Q ” is not arbitrary if both $j(P)/j(Q)$ and $j(Q)/j(P)$?

This idea has merit. It does not apply as a critique of the account presented here. If (i) “ P iff Q ” is a non-arbitrary prediction where $j(P)/j(Q)$; and (ii) it is not that P entails Q ; then (iii) there will always be some S whereby “ P iff S ” is logically equivalent to “ P iff Q ,” and where both $j(P)/j(S)$ and $j(S)/j(P)$. Proof: if (i) and (ii), let S be “ $P \ \& \ Q$.”

4.2 Shouldn't Non-Arbitrariness Hold Across Conjunction?

Consider:

Non-Arbitrary Conjunction: If hypothesis A and hypothesis B individually only make correct non-arbitrary predictions, then so should hypothesis “ $A \ \& \ B$ ”.

On the face of it, this seems a reasonable requirement. The account provided here violates it. To illustrate: you see a blue jay perched on a red mailbox. “All jays are blue” and “All mailboxes are red” both make non-arbitrary correct predictions. But “All jays are blue & all mailboxes are red” makes the arbitrary prediction “(There exists a blue jay or a mailbox) iff (there exists a jay or a red mailbox).” The mailbox evidence alone can justify the first proposition without justifying the second; the jay evidence alone can justify the second proposition without justifying the first.

Structurally, if two hypotheses make non-arbitrary predictions about *different* pieces of evidence, then a hypothesis that is their conjunction will make an arbitrary prediction. Because the different pieces of evidence are independently justified, the combined hypothesis, in a sense, claims too much in tacking the hypotheses together.

There are two obvious responses here. First, hold this as non-problematic. “All jays are blue” and “All mailboxes are red” would each be directly supported by a blue jay on a red mailbox, and “All jays are blue & all mailboxes are red” would not. The upside is clarity and simplicity, and the logic will carve out groupings of hypotheses based on the structure of the evidence. This has useful implications (Sections 4.3 and 4.4). The downside would be socio-disciplinary: the approach will likely face on-going demands to explain counter-intuitive disguised conjunction cases (e.g. Section 5.1).

The second response would be to simply lexically describe the different ways the hypotheses are supported. “All jays are blue” and “All mailboxes are red” would be directly supported by a blue jay on a red mailbox, and “All jays are blue & all mailboxes are red” would be indirectly supported as the conjunction of directly supported hypotheses. For expository clarity, this article focuses on which hypotheses make arbitrary predictions and which do not. It leaves open which of the above approaches is better.

The fact that a plausible hypothesis makes an arbitrary prediction need not necessarily render it inductively bad – if it is equivalent to the conjunction of hypotheses that individually only make correct and non-arbitrary influences.

4.3 The Explanatory Power of Conjunctive Relations

The issues of 4.2 arise when two hypotheses make non-arbitrary correct predictions, but their conjunction does not. A different scenario would be when H1 and H2 both make non-arbitrary correct predictions and “H1 & H2” does as well.

This case is interesting, not because there is any obvious problem – all three hypotheses pass the non-arbitrariness test. It is interesting because the conjunction hypothesis might be able to explain *why* the sub-hypotheses are supported, as the same feature of the evidence must support all three.

For example, consider evidence establishing the existence of a white swan and the hypothesis that “All gulls are white.” This hypothesis makes correct and non-arbitrary predictions.

To understand why, we can conjunct it with another hypothesis that does so, too, such as “All birds that are not gulls are white,” where the combined hypothesis is also non-arbitrary, in this case, “All birds are white.” Hence, “All gulls are white” is supported by a white swan, at least partly, *because* the white swan supports “All birds are white.”

This method is particularly useful for superficially puzzling cases. Take the white swan and the hypothesis “All robins are red.” This makes correct and non-arbitrary predictions.

The previous bird-colour type of conjunction is not an option here: “All robins are red & All birds that are not robins are red” makes incorrect predictions. But the following makes only correct non-arbitrary predictions: “All robins are red & All robins are not-red,” or, as it can be rephrased, “All birds are not-robins.”

This indicates that “All robins are red” is supported *because* a white swan supports the hypothesis that there are no robins. This makes sense – absent all other evidence, swans should support scepticism about non-swans, such as robins. If this still seems intuitively odd due to smuggled-in background evidence of robin prevalence, replace robins with dodos or phoenixes. If you keep observing birds that are not phoenixes, you should, absent other background evidence, become more sceptical that there are any phoenixes.

Interestingly, a particular single explanatory conjunction may not be the only explanation for why a hypothesis is supported. For instance, absent all other evidence, a white swan supports “All gulls are white” on two distinct grounds. The hypothesis is true if “All birds are white.” And it is true if “All birds are swans.” They are “distinct” as their conjunction yields arbitrary predictions, whereas they individually do not. In contrast, “All robins are red” is only supported via “All birds are swans.” Conjunctive exploration might be a valuable conceptual resource of the non-arbitrariness approach.

4.4 Entailment Analysis

The test for arbitrariness applies to the entailments of hypotheses. This has three immediate methodological consequences.

First, if hypothesis H makes a correct arbitrary prediction, then there is no X such that “H & X” renders the correct prediction non-arbitrary. There is nothing, for example, that can be tacked onto “All emeralds are grue” that will remove the arbitrary prediction about the green OT emerald.

Second, for any set of evidence, there will always be a unique set of hypotheses where both (i) each hypothesis makes only correct and non-arbitrary predictions and (ii) the conjunction of any two of the hypotheses makes an arbitrary prediction.

Call this the hypothesis “base-set” or each a “base-hypothesis.” Note: if conjunction did hold across non-arbitrariness, then there would always only be one member, and this section’s tool of base-set analysis would not exist.

Third, all hypotheses individually entailed by a member of the base-set, and logically compatible with all base-set members, will each make only correct and non-arbitrary predictions. All other logically possible hypotheses will not. One fertile implication is that, once the base-set is identified, it can be used to assess *why* a hypothesis is supported by the evidence, namely by looking specifically at which base-set member or members entail that hypothesis.

For example, in principle, evidence establishing the existence of a white swan will have a hypothesis base-set with two members: “All objects are white” and “All objects are swans.” “All objects are white swans” would make arbitrary predictions but, of course, is simply the conjunction of the other two.

In reality, in this situation, the typical evidential set will be heterogeneous and diverse. We will know many, many things about the swan concerning the observation, the location, ourselves, and much more. We will be awash in evidence. But, for illustrative purposes, focussing narrowly on the white swan evidence, it follows that any hypothesis entailed by either of the two base-set hypotheses, and logically compatible with the other, will make correct and non-arbitrary predictions about the white swan. Some examples: “All swans are white,” “All tomatoes are white,” “All shoes are yellow,” and “All parrots are non-white.”

Crucially, *which* base-hypothesis entails the derivative hypothesis in question will identify *why* the overall evidence supports it. “All parrots are non-white,” for example, is entailed by “All objects are swans.” Both gain their support from the evidence establishing that there is a swan. Whereas “All tomatoes are white” is also entailed by “All objects are white,” and gains its support in addition from evidence establishing that there is a white thing.

Some of these examples might seem intuitively odd. But that is down to either the implicit inclusion of currently not-included background evidence, or misconceptions about the truth conditions of “All F are G” hypotheses, as Section 5.3 sets out.

Two final intriguing features to note. First, as hypotheses can be entailed by more than one member of the base-set, there thus appears here to be a concept of relative support. If H1 is entailed by one member, and H2 by that member plus another, H2 is better supported.

Second, if the evidence is uncertain, that will translate into uncertainty between different base-sets and their entailments. As such, if we could quantify the evidential uncertainty, we should be able to quantify the relative uncertainty between hypotheses.

Stepping back, and keeping the previous discussion in mind, arbitrariness theory has three complementary ways of identifying when and why hypotheses are supported by evidence.

First, scrutinise the biconditional entailments of a given hypothesis on a case-by-case basis, as was initially done with the grue cases.

Second, see if puzzling cases can be explained by either (i) conjoining a non-arbitrary hypothesis with others while retaining the non-arbitrariness or (ii) splitting an arbitrary hypothesis into sub-hypotheses that lack the arbitrariness.

Third, ascertain which members of the hypothesis base-set of a body of evidence entail a given hypothesis. The member(s) in question help identify which specific features of the evidence cause the hypothesis to be supported.

As such, the tools can be used to assess why apparently puzzling hypotheses – such as “All unicorns are rainbow-pink” – are supported by diverse objects, such as a grey shoe, green emerald, or brown horse, yet storybooks with rainbow-pink unicorn illustrations also make unwarranted assumptions.

5 Case-Based Objections

Every vaguely plausible-looking grue resolution will address at least one type of grue case. And if a paper *only* discusses that, readers will be left having to puzzle through what the best interpretation of the condition implies for all other cases.

Furthermore, the risk here is literature ping-pong. A proposal addresses only one type of case; a response raising a problem case bats the debate back; the proposal is reformed or interpreted differently to address that case, and so on. Yet, at the same time, the set of hypotheses and objects is unbounded: it is neither possible nor desirable to crank through a massive list of possibilities.

Luckily, the list of core cases – those historically prominent or a problem for many suggestions – is relatively short. It should be straightforward for any proposal to state what is entailed for them. These seem to be of the following type:

Case	Hypothesis example	Object/observation	Challenge/right response?
Exclusionary Grue	All emeralds are [(green and OT) or (blue and not-OT)]	A green OT emerald	Should not support
Inclusionary Grue	All emeralds are [(green and exist on a Monday) or (blue and exist on a non-Monday)]	A green emerald existing on a Monday	Should not support
Emerose1	All (observed emeralds or unobserved roses) are green	A green observed emerald	Conventionally: should not support. Actually: fine
“ ”	“ ”	An observed non-green rose	Should undermine
Emerose2	All (emeralds or unobserved roses) are green	A green emerald	Conventionally: should not support. Actually: fine
“ ”	“ ”	An observed non-green rose	Should undermine
Operald	All (opals or unsampled emeralds) are blue	Sampled blue opals <i>and</i> sampled green emeralds	Collectively, should not support, should undermine

Each of these has been previously discussed explicitly, and this article's proposal delivers all the results in the right-hand column (the non-conventional ones where relevant). Readers may disagree with some results, at least though what those results are is clear. In addition, there may be further apparent problem cases that are proposal-specific. Below are some that are relevant to this article's approach.

5.1 Arbitrary Hypotheses and Disguised Conjunction Cases

“All swans are white and two-legged” will make arbitrary predictions about a white two-legged swan, such as (“There exists a white swan” iff “There exists a two-legged swan”). This seems problematic as it appears to be a straightforward, good inductive generalisation.

However, this is merely a slightly disguised conjunction case. The hypotheses “All swans are white” and “All swans are two-legged” both make correct and non-arbitrary predictions.

The earlier conjunction case was illustrated with evidence about two distinct objects: the jay and the mailbox. But, the underlying logic also applies to two distinctly evidenced properties of the same object.

5.2 Hypotheses and Evidence that are Judged Arbitrary where this Seems Problematic

Take the following three hypotheses and evidence establishing the subsequent proposition:

- “Any bird is pink iff it is a flamingo” and pink flamingos.
- “Only mammals produce milk” and mammals producing milk.
- “A pendulum makes one swing per second iff it is one meter long” and a one-meter pendulum making one swing per second.³⁵

³⁵ I'm grateful to an anonymous reviewer for suggesting the latter two as seemingly problematic for the approach.

Each of these would make arbitrary predictions about the evidence, such as, with the first, “there exists a pink bird iff it is a flamingo.”

This raises two immediate worries. One, intuitively the hypotheses make good inductions, and should not be judged arbitrary. Indeed, they seem exactly the sort of law-like generalisations prevalent in science. Two, if the evidence cited does not support them, then what would? We have apparently ruled out anything doing so, which seems dogmatic and highly restrictive.

However, this is too quick. The following hypotheses make only non-arbitrary correct predictions about the earlier evidence.

- “All flamingos are pink.”
- “All mammals produce milk.”
- “All 1m pendulums make one swing per second.”

Presumably, what makes “Any bird is pink iff it is a flamingo” persuasive over possible rivals, such as “All birds are pink,” is not just evidence of pink flamingos but evidence of non-flamingo birds being non-pink. But in that case, “All flamingos are pink” and “All non-flamingo birds are non-pink” would make correct non-arbitrary predictions, and this is a conjunction case. “All birds are pink” would now trivially make an incorrect prediction.

The same holds for the other hypotheses. Their plausibility *over potential rivals*, such as “All animals produce milk” and “All pendulums swing once per second,” hinges on evidence of non-milk-producing non-mammals and evidence of non-1m pendulums not swinging once per second.

In effect, given the cited evidence, the initial hypotheses claim too much. They are not bad hypotheses that should be ruled out, and they can be indirectly supported if that is an endorsed option. Instead, they gain their plausibility via background evidence, and that needs including to determine which of the rival hypotheses is comparatively better supported.

5.3 Hypotheses that are not Judged Arbitrary Where this Seems Initially Problematic

Take evidence establishing that there is a pink flamingo and the following hypotheses:

- F1: “All (flamingos or geese) are pink”
- F2: “All (flamingos or geese-sitting-down) are pink”
- F3: “All geese are pink”
- F4: “All geese are grey”

All these hypotheses make correct and non-arbitrary predictions. But, on the face of it, this appears to be the wrong result for a host of reasons. First, it seems a pink flamingo shouldn’t tell us anything about the colour of geese; it’s irrelevant. Second, we can substitute many other things for the geese-based clauses; thus, we have an apparent structural flaw where almost any generalisation can be smuggled in. Third, while a grey goose sitting down would contradict F1, F2 and F3, a grey goose standing up would only contradict F1 and F3. Yet, it seems that F2 should be undermined in this case.

The first two, however, rest on mistakes, and the third is straightforwardly handled. Namely, *absent other evidence*, a pink object should support the hypotheses “All objects are pink” and “All X’s are pink,” which is partly why we do learn things about non-flamingos, and partly why F1, F2 and F3 are supported. The earlier use of conjunction can demonstrate this: “F3 & all non-geese birds are pink” – that is “All birds are pink” – makes correct and non-arbitrary predictions. Or, the base-hypothesis explanatory method can be used: “All objects are pink” entails each of F1, F2 and F3.

However, that explanation seems to make F4 *more problematic*, as the geese in F4 are not pink. Yet, absent other evidence, a flamingo should support the hypothesis “All objects are flamingos,” and as such “All objects are not-Y” where Y is anything that is a non-flamingo. F4 is the hypothesis “All objects are not (non-grey geese).” The double negative makes quick cognitive parsing less easy, but the logic is straightforward.

Non-arbitrary tacking can clarify this: “F4 & all geese are non-grey” – that is “All birds are not geese “ – makes only correct non-arbitrary predictions. Or, the base-hypothesis explanatory method can be used: “All objects are flamingos” entails F4.

What misleads about these hypotheses in terms of initial intuitions is that we do have *other* evidence as to the colour of geese and that they exist. But, if we include that, then arbitrariness theory gets the right results: evidence establishing a grey goose was not sitting down would contradict F1 and F3, render F2 arbitrary, and support F4.

Some may also forget that “All F are G” has two mutually exclusive but individually sufficient truth conditions: (i) that there exist one or more Fs, and all are G; or (ii) that there are no Fs. Non-included background evidence of Fs may therefore cognitively mislead in cases where evidence supports the hypothesis only via (ii). This could suggest – wrongly and perhaps perplexingly – that the evidence supports greater confidence in (i).

The flamingo cases above are disguised and blended versions of emerald-like and ravens-paradox-like cases. The correct result concerning the specific cited evidence or object can appear wrong if background evidence implicitly drives the intuitive judgements or if we cognitively respond to the wrong features of the hypothesis or evidence. Once such background facts are explicitly included, and the conjunctive or base-hypothesis methods are used to identify the relevant features, the cases become clear.

Superficially problematic cases as such must be scrutinised on at least three grounds. First, if a hypothesis makes an arbitrary prediction about two distinct features or pieces of evidence, is it equivalent to the conjunction of hypotheses individually making correct and non-arbitrary predictions? Second, if a hypothesis seems inductively good but is arbitrary, does the judgement of its inductive goodness implicitly rely on other evidence? The test would be whether including the other evidence means the hypothesis is now supported (as a conjunction) over potential rivals. Third, given a particular piece of evidence, if a hypothesis seems irrelevant or arbitrary but is not ruled out, are we reacting to the wrong features of the evidence or hypothesis? The conjunctive or base hypothesis methods can be used to identify *why* the hypothesis makes non-arbitrary correct predictions. We don't just need to always rely upon intuitions. There is a logic to non-arbitrariness.

6 Conclusions

This article has explored four core possibilities. First, we can use the logical relations between the sets of justifications of premises to test whether bi-conditional (“P iff Q”) predictions are arbitrary. Second, doing so can capture that, for some grue-like G's, “All F are G” will make correct and yet also arbitrary predictions about the observation of an F that is G. Third, the idea of testing entailments for non-arbitrariness seems to have some structural advantages over predicate-generalisation rivals. Fourth, the approach may have built-in explanatory tools for identifying *why* predictions are good or bad, offering a different heuristic than mere case-based intuition endorsement.

Acknowledgments: I am grateful to Stephen Butterfill for his comments on a previous version.

Funding information: Author states no funding involved.

Author contribution: The author confirms the sole responsibility for the conception of the study, presented results and manuscript preparation.

Conflict of interest: Author states no conflict of interest.

References

- Achinstein, Peter and Stephen Barker. “On the New Riddle of Induction.” *Philosophical Review* 69 (1960), 511–22.
- Armstrong, David M. *What is a Law of Nature?* Cambridge: Cambridge University Press, 1985.
- Aune, Bruce. *Knowledge, Mind and Nature*. New York: Random House, 1967.

- Bealer, George. *Quality and Concept*. Oxford: Oxford University Press, 1982.
- Chisholm, Roderick M. *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press, 1957.
- Coakley, Mathew. *Motivation Ethics*. London: Bloomsbury Academic, 2017.
- Coakley, Mathew. "Hempel, Grue and the Logical Empiricist Baseline." *Erkenntnis* 83 (2018), 969–82.
- Dorst, Chris. "Bet Accepted: A Reply to Freitag." *Philosophical Quarterly* 68 (2018), 175–83.
- Freitag, Wolfgang. "I'll Bet You Solve Goodman's Riddle." *Philosophical Quarterly* 65 (2015), 254–67.
- Freitag, Wolfgang. "The Disjunctive Riddle and the Grue-Paradox." *Dialectica* 70 (2016), 185–200.
- Freitag, Wolfgang. "Why Doxastic Dependence Defeats Grue: A Response to Dorst's Reply." *Philosophical Quarterly* 69 (2019), 156–65.
- Godfrey-Smith, Peter. "Goodman's Problem and Scientific Methodology." *The Journal of Philosophy* 100 (2003), 573–90.
- Godfrey-Smith, Peter. "Induction, Samples, and Kinds." In *Carving Nature at Its Joints: Topics in Contemporary Philosophy*, edited by Joseph K. Campbell, Michael O'Rourke, and Matthew H. Slater, 33–52. Cambridge, MA: MIT Press, 2011.
- Goodman, Nelson. "A Query on Confirmation." *Journal of Philosophy* 43 (1946): 383–85.
- Goodman, Nelson. *Fact, Fiction and Forecast*. 4th ed. Cambridge: Cambridge University Press, 1983.
- Goodman, Nelson. *Problems and Projects*. Indianapolis: Bobbs-Merrill, 1972.
- Jackson, Frank. "Grue." *Journal of Philosophy* 72 (1975), 113–31.
- Jackson, Frank and Robert Pargetter. "Confirmation and the Nomological." *Canadian Journal of Philosophy* 10 (1980), 415–28.
- Moreland, James. "On Projecting Grue." *Philosophy of Science* 43 (1976), 363–77.
- Okasha, Samir. "What Does Goodman's 'Grue' Problem Really Show?" *Philosophical Papers* 36, no. 3 (2007), 483–502.
- Quine, Willard Van Orman. "Natural Kinds." In *Essays in Honor of Carl G. Hempel*, edited by Nicholas Rescher et al., 41–56. Dordrecht: D. Reidal Publishing, 1970.
- Roskies, Adina. "Robustness and the New Riddle Revived." *Ratio* 21 (2008), 218–30.
- Schramm, Alfred. "Evidence, Hypothesis, Grue." *Erkenntnis* 79 (2014), 571–91.
- Sellars, Wilfrid. *Science and Metaphysics: Variations on Kantian Themes*. Atascadero, CA: Ridgeview Publishing, 1967.
- Shoemaker, Sydney. "Properties, Causation and Projectibility." In *Applications of Inductive Logic*, edited by L. Jonathan Cohen and Mary Hesse, 291–312. Oxford: Oxford University Press, 1980.
- Stalker, Douglas, ed. *Grue! The New Riddle of Induction*. Chicago: Open Court, 1994.
- Warren, Jared. "The Independence Solution to Grue." *Philosophical Studies* 180 (2023), 1305–26.
- Wilkerson, T. E. "A Gruesome Note." *Mind* 82 (1973), 276–77.