

A Phenomenological Approach to the Bayesian Grue Problem

Ibrahim Dagher*

University of California, Davis

It is a common intuition in scientific practice that positive instances confirm. This confirmation, at least purely based on syntactic considerations, is what Nelson Goodman's 'Grue Problem', and more generally the 'New Riddle' of Induction, attempt to defeat. One treatment of the Grue Problem has been made along Bayesian lines, wherein the riddle reduces to a question of probability assignments. In this paper, I consider this so-called Bayesian Grue Problem and evaluate how one might proffer a solution to this problem utilizing what I call a phenomenological approach. I argue that this approach to the problem can be successful on the Bayesian framework.

1 Introduction

It is a common intuition in scientific practice that positive instances confirm. That is, that repeated instantiations of some predicate P lend inductive support to a general hypothesis wherein P is projected. The hope that such syntactic considerations might serve as the basis for an inductive logic is what Goodman¹ sets out to defeat in his so called 'New Riddle' of induction. As such, the New Riddle has received considerable discussion, including treatments of the riddle along Bayesian lines.² One such reformulation has been proffered by Sober (1994), which prompts new considerations — such as how different *kinds* of hypotheses differ with respect to their confirmation conditions — and how this might give rise to various manifestations of the riddle. In this paper, I consider the New Riddle cast in the Bayesian framework proposed by Sober, and appraise a 'phenomenological approach' to the riddle. I argue that the approach, as applied to the grue problem, can be successful. I will proceed as follows. In §2, I explicate a Bayesian formulation of the grue problem along the lines Sober³ outlines. In §3, I discuss some general difficulties Bayesian answers will have to deal with. I outline the phenomenological approach to answering the problem in §4 before concluding in §5.

2 A Bayesian Grue Problem

Consider the predicate 'grue', which applies to any x just in case it is green and examined earlier than some time t or blue and examined at or later than t . Following Sober (1994), we can now begin to concern ourselves with various hypotheses from which the riddle will emerge. First, consider these two hypotheses, which are said to be *generalizations*:

*My name is Ibrahim Dagher, and I'm a second-year philosophy student at UC Davis. I have a special research interest in epistemology, as well as philosophy of law and religion. I enjoy playing basketball and doing the daily Wordle!

1. Nelson Goodman, *Fact, Fiction, and Forecast* (Harvard University Press, 1955).
2. For formulations other than Sober's, see Irving John Good, "Explicativity, Corroboration, and the Relative Odds of Hypotheses," *Synthese* 30, nos. 1-2 (1975): 39-73; Richard C. Jeffrey, *The Logic of Decision* (New York, NY, USA: University of Chicago Press, 1965)
3. Elliott Sober, "Grue!: The New Riddle of Induction," chap. No Model, No Inference: A Bayesian Primer on the Grue Problem, ed. Douglas Stalker (Open Court Publishing Group, 1994), 225-238.

1. **AllGreen:** All emeralds are green.
2. **AllGrue:** All emeralds are grue.

Presumably, AllGreen is a perfectly rational generalization to commit oneself to. However, AllGrue does not seem to be. Thus, the first question of the riddle is this: what asymmetry exists between AllGreen and AllGrue, such that we are justified in our belief in the former *rather* than the latter?

There is also another question to be asked at this point. Consider these two hypotheses, which are instead said to be *predictions*:

1. **NextGreen:** The next emerald to be examined will be green.
2. **NextGrue:** The next emerald to be examined will be grue.

Again, presumably it would only be rational to believe the first prediction, assuming the next emerald will be examined at or later than t , so that these predictions are contradictory. We are thus compelled to ask: what asymmetry exists between NextGreen and NextGrue, such that we are justified in our belief in the former rather than the latter? There are two distinct issues at hand: the first is finding some epistemic asymmetry between AllGreen and AllGrue, and the second is finding one between NextGreen and NextGrue.

Finally, I wish to make one more distinction. The specific hypotheses and their respective questions, as formulated herein, are what I take to constitute the grue problem — the problem of finding some epistemic asymmetry between AllGreen and AllGrue, and NextGreen and NextGrue *specifically*. The New Riddle is the problem of characterizing the epistemic relationships between hypotheses of generalizations, predictions, and their respective instantiations *more generally*. This distinction is important because the solution I propose here ought to be considered only a solution to the grue problem, and not the much more general riddle.

2.1 Bayesian Confirmation Conditions

With these questions on the table, we can now move to explaining what the sufficient conditions are for answering these questions. As mentioned above, these conditions shall be cast along Bayesian lines.

First, since the sought conclusion to both of our questions will take the form 'hypothesis H_1 can be assigned a higher *posterior probability* than hypothesis H_2 because...' it is worth explaining what obtaining a posterior probability of a given hypothesis consists of for Bayesians. Where H abbreviates some hypothesis and O abbreviates the set of observations we have made, Bayes' theorem tells us that the posterior probability of H can be calculated by reference to the likelihood and prior probability of H , as well as the probability of O :

$$\Pr(H | O) = [\Pr(O | H) * \Pr(H)] / \Pr(O)$$

The $\Pr(H | O)$ is the posterior probability of H —the probability that H is true given the observations we have made. $\Pr(O | H)$, on the other hand, is the *likelihood* of H : the probability H confers onto O 's obtaining. Lastly, the $\Pr(H)$ is what is often termed the *prior* probability of the hypothesis: the probability H enjoys before any observations are made.

Since the nature of both of our questions is comparative, we should wish to reformulate Bayes' theorem into a comparative principle. This is simple enough:

(CPs): $\Pr(H_1 | O) > \Pr(H_2 | O)$ when and only when

$$[\Pr(O | H_1) * \Pr(H_1)] > [\Pr(O | H_2) * \Pr(H_2)]$$

Interestingly, the comparative principle as explicated above is a *synchronic* one. We might also wonder what difference in the probability of H is *incited* by the truth of O. In other words, we may also be interested in a *diachronic* comparative principle. Assuming that the larger the difference between the posterior and prior probabilities of a hypothesis, the greater the confirmation, then:

(CP_d): O confirms H₁ more than H₂ when and only when

$$[\Pr(H_1 | O) - \Pr(H_1)] > [\Pr(H_2 | O) - \Pr(H_2)]$$

Thus, (CP_d) differs from (CP_s). So, two further subdivisions have to be made with respect to the issues at hand: not only must we consider the relevant probabilities of AllGreen compared to AllGrue and NextGreen compared to NextGrue, but each comparison must be considered diachronically and synchronically. Let us diagnose each in turn.

2.2 AllGreen vs AllGrue: A Synchronic Analysis

Suppose 'D' denotes a proposition that contains the relevant past data, namely, 'all emeralds examined have been observed to be green'. To analyze the posterior probabilities of AllGreen and AllGrue synchronically, it is important to begin with what is commonly affirmed: the truth of either AllGreen or AllGrue entails D. Thus, the likelihoods of either hypothesis are exactly 1. This is just the fact that our past data confirms both generalizations.

However, given (CP_s), if two hypotheses are of equivalent likelihoods, the only way in which one could have a higher posterior probability than the other is if one has a higher prior probability than the other. That is:

$$\Pr(\text{AllGreen} | D) > \Pr(\text{AllGrue} | D) \text{ when and only when} \\ \Pr(\text{AllGreen}) > \Pr(\text{AllGrue})$$

If this is correct, our condition for preferring AllGreen rather than AllGrue is this: AllGreen enjoys a higher prior probability than AllGrue. More will have to be said about what might qualify — or if anything at all can qualify — as a justified reason for such prior probability assignments.

2.3 AllGreen vs AllGrue: A Diachronic Analysis

Much like the synchronic analysis, the conditions for different posterior probabilities on the diachronic analysis appear to reduce to considerations of prior probability. Using CP_d with Bayes' theorem, we obtain the following for AllGreen:

$$[[\Pr(D | \text{AllGreen}) * \Pr(\text{AllGreen})] / \Pr(D)] - \Pr(\text{AllGreen})$$

And the same for AllGrue:

$$[[\Pr(D | \text{AllGrue}) * \Pr(\text{AllGrue})] / \Pr(D)] - \Pr(\text{AllGrue})$$

As an inequality, this transforms into:

$$[[1 - \Pr(D)] * \Pr(\text{AllGreen}) / \Pr(D)] > [[1 - \Pr(D)] * \Pr(\text{AllGrue}) / \Pr(D)] \text{ And, on the assumption that } \Pr(D) < 1, \text{ we obtain:} \\ \Pr(\text{AllGreen}) > \Pr(\text{AllGrue})$$

Thus, on our diachronic analysis the posterior probabilities are higher for the AllGreen hypothesis than the AllGrue hypothesis when and only when the priors are higher *and* our data was not certain.

2.4 NextGreen vs NextGrue: A Diachronic Analysis

For our predictive hypotheses, the conditions under which we can assign comparatively higher posterior probabilities change. I will begin with the diachronic case. At first glance, it might be thought that because our past data confirms and raises the probability of the general hypotheses, it ought to also confirm and raise the probability of the predictive hypotheses. After all, the truth of either general hypothesis *entails* the truth of the respective predictive hypothesis.

However, this is not so. At least, not without significant assumptions about the sampling process involved. If I know that some marbles placed in a bag were randomly sampled from a source with an equivalent ratio of black to blue to red to green marbles, then the fact that every marble I have examined has been red does not confer any further probability on the predictive hypothesis ‘the next marble will be red’. The probability remains 0.25. Yet, the fact that every marble I have examined has been red *does* increase the probability that every marble is red, by virtue of the fact that this has eliminated certain hypotheses from the possibility space (namely, all the hypotheses entailing that less than x-many red marbles would be examined, such as the hypothesis that all the marbles are black).

This asymmetry in confirmation arises precisely because of my knowledge of the sampling process. My knowing that the marbles do not have their colors selected, as it were, *collectively*, or by some law-like process, precludes the possibility that all the marbles’ being homogeneous in color is anything other than mere happenstance. It is only when this possibility is introduced that one can begin to alter the probability of a predictive hypothesis.⁴

In other words, it is only when the conjunction of the relevant prediction and the data is more probable than the independent occurrence of each that the data confirms the prediction. Evidence confirms a prediction only if the two are *positively correlated*, or dependent, facts. If they are independent, then their conjunction can never be more probable than the occurrence of both of their conjuncts.

With this analysis in hand, we are now prepared to outline the probabilistic conditions on which we ought to prefer NextGreen over NextGrue. First, assume that the next emerald observed will be either green or blue. Next, assume the present moment is *t*, so that NextGreen and NextGrue are contradictory, and logically exhaustive, hypotheses. So, if some condition confirms NextGreen, it will disconfirm NextGrue. Here is the condition:

(C_d): NextGreen is confirmed by data D if and only if

$$\Pr(\text{NextGreen} \ \& \ D) > [\Pr(\text{NextGreen}) * \Pr(D)]$$

Why might we think that the probability of the conjunction of NextGreen and our past data is greater than the independent occurrence of each of these facts? Presumably it is because of an assumption about the nature of emeralds and their color: namely, the color predicate that is ultimately true of emeralds should be true of them *qua* their being emeralds. That is, we assume their color is determined as a *group*. It is not as though each emerald is sampled from a possible space of colors individually and independently of any other emerald. The more pressing question that arises at this point is not that of why we might think the inequality would hold, but rather why we should think *this* inequality holds. Plausibly, NextGrue is also positively associated with the past data in the same way that NextGreen is. Our motivations for thinking that emeralds would collectively be green apply equally well for thinking that emeralds would collectively be grue.

This question will soon be addressed, but the important lesson here is this: NextGreen and NextGrue have slightly different conditions for epistemic asymmetry than do AllGreen and AllGrue. For our past data to confirm the generalizations, we need some reason to prefer a certain assignment of priors. For our past data to confirm the predictions, we need some reason to prefer a certain positive association over another.

4. For more on this relationship, see Sober/Elliott Sober, “Confirmation and Law-Likeness,” *Philosophical Review* 97, no. 1 (1988): 93–98

2.5 NextGreen vs NextGrue: A Synchronic Analysis

Finally, let us consider how our past data might serve to confirm the predictive hypotheses on a synchronic analysis. Holding fixed the aforementioned conditions that made it such that NextGreen and NextGrue were contradictory and logically exhaustive hypotheses, on a synchronic analysis the question of posterior probability assignment boils down to the following:

When is $\Pr(\text{NextGreen} \mid D) > \Pr(\text{NextGrue} \mid D)$?

Since $\Pr(\text{NextGrue}) = 1 - \Pr(\text{NextGreen})$, this can be expanded to:

$$\Pr(\text{NextGreen} \ \& \ D) - [\Pr(\text{NextGreen}) * \Pr(D)] > \\ [\Pr(D) * [1 - 2 * P(\text{NextGreen})]] / 2$$

Simplifying:

$$\Pr(\text{NextGreen} \ \& \ D) > [\Pr(D) / 2]$$

This becomes:

$$\Pr(\text{NextGreen} \mid D) > 0.5$$

Thus, the synchronic case is similar to the diachronic case: insofar as NextGreen is positively associated with D, and the $\Pr(\text{NextGreen} \mid D) > 0.5$, it follows that $\Pr(\text{NextGreen} \mid D) > \Pr(\text{NextGrue} \mid D)$.

Unsurprisingly, there remains a kind of epistemological indeterminacy under both analyses with respect to which predictive hypothesis ought to be assigned the aforementioned positive association. There is nothing, it appears, in D that could possibly account for that kind of preferential assignment.

3 Difficulties on the Bayesian Framework

Very roughly, answers to the Bayesian grue problem will have to amount to a favorable prior probability assignment to AllGreen rather than AllGrue, and a favorable assignment of positive association between NextGreen and the past data rather than NextGrue. But before attempting to characterize a certain answer as meeting these conditions, there appear to be deeper difficulties with even meeting these conditions at all.

The first difficulty is the well-known *problem of the priors*. That is, what norms ought to dictate the distribution of prior probabilities to any logically exhaustive set of hypotheses under consideration? Is the only requisite norm a requirement on cohering with the axioms of probability (Subjective Bayesianism)? Or is there, in addition to this norm, a norm on which our priors follow some concern for evidential or reason-based indifference (Objective Bayesianism)? Or is it rather that our priors should be such that conditionalizing on a given class of evidence produces posteriors that would be in line with some theoretical virtue like explanatory power, simplicity, or convergence to truth (Future-Oriented Bayesianism)? This is an ongoing debate, and if an answer to the grue problem requires that our past data supports AllGreen *rather than* AllGrue if and only if the prior of the former is greater than the prior of the latter, any attempted solution will have to contend with the question of what kinds of norms ought to dictate our prior probability assignments.

Sober (1994) raises another obstacle for any solution to the Bayesian grue problem. It can be put as follows: either the prior probabilities we are considering are objective or subjective. If they are objective, then they cannot possibly be assigned *a priori*. If they are subjective, then varying prior probability assignments could not possibly amount to an epistemic asymmetry between AllGreen and AllGrue.

However, this dilemma, at least with respect to the grue problem, can be dissolved. It is obvious that the probabilities discussed with respect to AllGreen and AllGrue are not objective chances. Indeed, such a notion appears to be fraught in the context of the question of what colors emeralds instantiate. Rather, I take the probabilities discussed herein to be *credences*. The $\text{Pr}(\text{AllGreen} \mid D)$ represents the credence, or degree of belief, one has in the hypothesis AllGreen on the past data. However, *pace* Sober, this fact does not remove the possibility of an epistemic asymmetry. This would only be the case if our answer to the question ‘Why believe AllGreen *rather than* AllGrue?’ were the descriptive answer ‘Because *in fact* my credence in AllGreen given the data is greater than my credence in AllGrue given the data.’ But, as I see it, the answer we will take on is actually normative. It is of the form, ‘Because there is a (true) norm *N* according to which I ought to have a greater credence in AllGreen given the data rather than AllGrue given the data.’ Surely if I ought to have a greater credence in some hypothesis compared to another hypothesis that constitutes a substantive epistemic asymmetry between the two hypotheses.⁵

4 The Phenomenological Approach

Having now remarked on what conditions the Bayesian grue problem requires of answers, and recognizing some general difficulties with any answer, I now wish to briefly sketch a solution that I call the ‘phenomenological approach’, and discuss how it provides new insights into the conditions and difficulties analyzed above.

Let’s begin with noting the following difference between AllGreen and AllGrue. If all emeralds are grue, and there are emeralds examined earlier and later than *t*, then there is a phenomenological asymmetry in the world. That is, before *t* we will have a certain phenomenological experience associated with the observation of emeralds. Then, after *t*, we will have a noticeably different phenomenological experience when we view emeralds. We might put the point as follows: grue emeralds are perceptually different.

On the other hand, if all emeralds are green, then there is phenomenological symmetry. That is, there is a constant phenomenological experience associated with the observation of emeralds, since green emeralds are perceptually the same.

The phenomenological approach to the grue problem attempts to draw an epistemic asymmetry between AllGreen and AllGrue on the basis of this difference.⁶ It is important, however, to recognize what is *not* being claimed as a difference between the two hypotheses. It is not being claimed that grue emeralds do not instantiate the same color across times (or some equivalently grue-ified predicate, ‘grulor’). Nor is the claim that grue emeralds *qua* emeralds experience some kind of change in their properties across times. The claim is merely relative to green speaking humans: if there are emeralds examined earlier and later than *t*, and these emeralds are grue, there will be a change in our phenomenological experience.

4.1 Building an Asymmetry

I take the lesson of the New Riddle to be the following: observing that a certain predicate *P* has consistently applied in the past does not *by itself* warrant the projection of *P* into the future. This is because there are many predicates *P*, *P*’...

5. I recognise that there are many pressing issues concerning the objectivity of norms, as well as the epistemology associated with being justified in their assertion. Unfortunately, because of space, I am unable to provide evaluations of these questions. What I am concerned with is whether there could be *any such norm* applicable to the grue problem.

6. For another attempt at utilizing this fact to solve the grue problem, see Sydney Shoemaker, ‘Functionalism and Qualia,’ *Philosophical Studies* 27, no. May (1975): 291–315, <https://doi.org/10.1007/BF01225748>.

that apply equally well of the same past phenomena, but if they were to be projected about phenomena examined in the future, would contradict P. There must be some other consideration, apart from consistent application to the past, that distinguishes P from P', P"...

It's at this point that the phenomenological approach begins. It notes that rather than projecting any predicates that have been true in the past, we ought to instead project the phenomenological experiences that have been had in the past.

Certainly, the principle that we ought to project the phenomenological experiences associated with past phenomena is just as, if not more, intuitive than the principle that we ought to project the predicates of past phenomena.

But now our epistemic asymmetry emerges. For while it is true that a whole host of predicates apply to the same past data, and thus the principle that 'predicates of past phenomena ought to be projected' is false, there is only *one* way in which we have experienced past data. There is only one phenomenological experience associated with the observation of emeralds, and thus the principle 'the phenomenological experience associated with the past data ought to be projected' cannot be defeated by consideration of the various predicates that might all equally apply to the data. And, since only AllGreen is consistent with the application of this principle, it is on this basis that we ought to prefer AllGreen to AllGrue.⁷

Another way of putting the idea is as follows. Goodman's New Riddle tells us that the following inductive schema (alone) is faulty:

1. a_1 is green

2. a_2 is green

3. a_3 is green

.

.

.

C: All a 's are green.

The phenomenological approach does not attempt to provide a semantic or epistemic explanation as to why only 'green', and not 'grue', will fit this schema in a truth preserving way. Rather, it proposes the alternative schema, Sp, with the principle that the only predicates that ought to be projected are those that entail a projection of perception:

1. a_1 is green and there is a single phenomenological experience p associated with observing a_1

2. a_2 is green and p is associated with observing a_2

3. a_3 is green and p is associated with observing a_3

.

.

.

Principle: If all a 's are green, then p is associated with observing all a 's C: All a 's are green

7. Interestingly, Barry Ward, "Explanation and the New Riddle of Induction," *Philosophical Quarterly* 62, no. 247 (2012): 365–385 uses a similar principle, built instead along explanatory rather than phenomenological lines, to reach this conclusion.

The idea at the heart of S_p is that a predicate ought to be projected for some a 's just in case (i) there are many positive instances of the predicate amongst the a 's and (ii) if the predicate were to be projected, that would entail projecting the same phenomenological experience that has been true of the past a 's.

I mentioned that the first inductive schema suffers from the following objection (which is just the grue problem): a substitution of 'grue' for 'green' in the argument yields a set of true premises, but a (supposedly) false conclusion. This schema alone, then, cannot be all there is to the inductive logic. Some semantic or epistemic considerations of the predicates inserted into the schema are at play, otherwise the schema is not truth preserving.

Might S_p suffer the same fate? Consider substituting 'grue' for 'green'. For each grue emerald it is true that the observation of that emerald has a phenomenological experience associated with it. So premises (1-...) are true. But the last premise, the principle, would be false. For if all emeralds are grue, then certainly p is not associated with observing all emeralds. In fact, if all emeralds are grue, then the p we have associated with all the past emeralds will not be the same for the emeralds examined later than t . The principle is built into the schema to discriminate between predicates whose projections entail different perceptual experiences and those that do not.

4.2 Bayesian Application

With the basic thrust of the phenomenological approach in mind, let us now turn to applying the idea in the context of the four Bayesian subdivisions of the grue problem.

4.2.1 AllGreen and AllGrue: Synchronic and Diachronic

Our Bayesian analysis in §2 took for granted that the likelihoods of AllGreen and AllGrue were the same. This was because both hypotheses entailed the relevant data proposition, D , which was 'all emeralds examined have been observed to be green'. It becomes immediately obvious that, on the phenomenological approach, this is *not* the relevant data proposition. The entire thrust of the solution proposed herein is that the relevant facts to our induction are not that x -many emeralds are green, but rather that x -many emeralds have a phenomenological experience associated with their being green.

So, the data we are considering must be expanded to include the phenomenological facts associated with our observations *and* the fact that the principle in our schema is only true with respect to AllGreen.⁸ Call this data proposition D' , which precisely spelled out is just this: 'premises (1-...) and the principle of S_p are true for the predicate 'green' and false for 'grue'".

Thus we are now interested in $\Pr(\text{AllGreen} \mid D')$ and $\Pr(\text{AllGrue} \mid D')$. As before, by Bayes' theorem this gives us:

$$\Pr(\text{AllGreen} \mid D') = [\Pr(D' \mid \text{AllGreen}) * \Pr(\text{AllGreen})] / \Pr(D')$$

$$\Pr(\text{AllGrue} \mid D') = [\Pr(D' \mid \text{AllGrue}) * \Pr(\text{AllGrue})] / \Pr(D')$$

I have already given some reason to think that our phenomenological data, D' , confers a higher credence to AllGreen rather than AllGrue. But note why this is the case: it is because the *likelihoods* of the hypotheses are now different, and not the priors. Why think that $\Pr(D' \mid \text{AllGreen}) > \Pr(D' \mid \text{AllGrue})$? First, note that:

$$\Pr(D' \mid \text{AllGreen}) = \Pr(D' \ \& \ \text{AllGreen}) / \Pr(\text{AllGreen})$$

8. It might be contended here that the principle I constructed in the schema, which constitutes the asymmetry between green and grue, is not part of our *data*, but instead should be considered a norm that governs prior probability assignment. I think this is mistaken. It does not follow from any analytic analysis of 'green' or 'grue' that one predicate should be true or false with respect to the principle. It seems possible that our perceptual experience of the world could have been such that we view grue emeralds as the same, and green emeralds as different. That our phenomenology happens to be one way rather than another is a fact we learn *a posteriori* and should be considered part of empirical data.

$$\Pr(D' \mid \text{AllGrue}) = \Pr(D' \ \& \ \text{AllGrue}) / \Pr(\text{AllGrue})$$

Assuming identical priors, the question becomes: why think that $\Pr(D' \ \& \ \text{AllGreen}) > \Pr(D' \ \& \ \text{AllGrue})$? Here is one argument: D' , which is just the truth of the premises of S_p with respect to 'green' and their falsity with respect to 'grue', provides a good basis for inferring AllGreen, but not AllGrue. And since the conjunction of a set of premises and a conclusion that can be justifiably inferred from the premises is more likely than a conjunction of those premises with some arbitrary conclusion (of the same prior) that cannot be justifiably inferred, it follows $\Pr(D' \ \& \ \text{AllGreen}) > \Pr(D' \ \& \ \text{AllGrue})$. Indeed, if we have a seemingly justified schema from D' to AllGreen, but not AllGrue, the likelihood of AllGreen will be greater than AllGrue.

Much the same can be said for the diachronic comparison of the two generalizations. Assuming that $\Pr(D') < 1$, by the same argument it follows that AllGreen has a higher likelihood, and therefore posterior probability, than AllGrue.

4.2.2 NextGreen and NextGrue: Diachronic and Synchronic

What about NextGreen and NextGrue? Starting with the diachronic condition:

(C_d): NextGreen is confirmed by data D' if and only if

$$\Pr(\text{NextGreen} \ \& \ D') > [\Pr(\text{NextGreen}) * \Pr(D')]$$

The central issue that arose was: why think that this inequality holds *rather than* the inequality with NextGrue? There was nothing in our previous data proposition that appeared to break this symmetry. However, utilizing D' , we now have a reason to think that only the above inequality holds. This is because the probabilities of the conjunctions can be reduced to a question of posterior probabilities:

$$\Pr(\text{NextGreen} \ \& \ D') = [\Pr(\text{NextGreen} \mid D') * \Pr(D')]$$

$$\Pr(\text{NextGrue} \ \& \ D') = [\Pr(\text{NextGrue} \mid D') * \Pr(D')]$$

And, much like earlier, there is good reason to think $\Pr(\text{NextGreen} \mid D') > \Pr(\text{NextGrue} \mid D')$. Namely, our phenomenological data confers a higher probability on the next emerald being green than it does for the next emerald being grue.

The same can be said in the synchronic case. D' gives us good reason to think that it is more likely that the next emerald is green rather than not. Note also that the arguments I gave in favor of a higher posterior for AllGreen rather than AllGrue, if successful, can provide reason for thinking the data confirms NextGreen over NextGrue, provided we assume that there is positive association between emerald colors.

5 Concluding Remarks

I have outlined how one might apply a so-called phenomenological approach to the grue problem construed in a Bayesian framework. I would like to now — even more briefly — remark on an interesting feature of this approach.

The interesting lesson seems to be this: the answer provided herein does not rely on justifying a higher prior to AllGreen or AllGrue as independent hypotheses. This is normally the principal difficulty with Bayesian treatments of the grue problem: if both predicates are interdefinable, on what basis might prior probabilities be assigned asymmetrically? By arguing instead that the data be expanded to include what phenomenological experiences applied in the past, all that

had to be justified was a different prior *relationship* the data bore to each hypothesis. Solutions to the grue problem should not aim to locate an asymmetry in $\Pr(\text{AllGreen})$ and $\Pr(\text{AllGrue})$ at all, but rather in the different likelihoods produced by some expanded data, D^* .⁹

I have construed the expanded data along phenomenological lines, but this need not be the case. Whatever asymmetry one finds between AllGreen and AllGrue to be the more pressing concern — be it phenomenological or other — ought to be construed as creating different likelihoods for the hypotheses by virtue of some expanded set of data to consider. Different assignments of priors are unnecessary. Reconsidering our data well appears to be enough.

References

- Good, Irving John. "Explicativity, Corroboration, and the Relative Odds of Hypotheses." *Synthese* 30, nos. 1-2 (1975): 39–73.
- Goodman, Nelson. *Fact, Fiction, and Forecast*. Harvard University Press, 1955.
- Jeffrey, Richard C. *The Logic of Decision*. New York, NY, USA: University of Chicago Press, 1965.
- Shoemaker, Sydney. "Functionalism and Qualia." *Philosophical Studies* 27, no. May (1975): 291–315. <https://doi.org/10.1007/BF01225748>.
- Sober, Elliott. "Confirmation and Law-Likeness." *Philosophical Review* 97, no. 1 (1988): 93–98.
- . "Grue!: The New Riddle of Induction." Chap. No Model, No Inference: A Bayesian Primer on the Grue Problem, edited by Douglas Stalker, 225–238. Open Court Publishing Group, 1994.
- Ward, Barry. "Explanation and the New Riddle of Induction." *Philosophical Quarterly* 62, no. 247 (2012): 365–385.

9. A similarly spirited remark is found in Fitelson (2008), wherein it is suggested that solutions to the grue problem can make use of the fact that 'the evidence only confirms both hypotheses *depending on the background corpus one starts with*'(emphasis added).