http://plato.stanford.edu/entries/learning-formal/

(On Goodman’s New Riddle of Induction)

This illustrates how means-ends analysis can evaluate methods: the bold method meets the goal of reliably arriving at the right answer, whereas the skeptical method does not. Note the character of this argument against the skeptic: The problem, in this view, is not that the skeptic violates some canon of rationality, or fails to appreciate the "uniformity of nature". The learning-theoretic analysis concedes to the skeptic that no matter how many black ravens have been observed in the past, the next one could be white. The issue is that if all observed ravens are indeed black, then the skeptic *never answers* the question "are all ravens black?". Getting the right answer to that question requires generalizing from the evidence *even though* the generalization could be wrong.

As for the bold method, it's important to be clear on what it does and does not achieve. The method will eventually settle on the right answer -- but it (or we) may never *be certain* that it has done so. As William James put it, "no bell tolls" when science has found the right answer. We are certain that the method will eventually settle on the right answer; but we may never be certain that the current answer is the right one. This is a subtle point. The next example illustrates this point further.

A Riddle of Induction

Nelson Goodman posed a famous puzzle about inductive inference known as the (New) Riddle of Induction ([Goodman 1983]). Our next example is inspired by his puzzle. Goodman considered generalizations about emeralds, involving the familiar colours of green and blue, as well as certain unusual ones:

> Suppose that all emeralds examined before a certain time *t* are green ... Our evidence statements assert that emerald *a* is green, that emerald *b* is green, and so on...
>
> Now let us introduce another predicate less familiar than "green’’. It is the predicate "grue’’ and it applies to all things examined before *t* just in case they are green but to other things just in case they are blue. Then at time *t* we have, for each evidence statement asserting that a given emerald is green, a parallel evidence statement asserting that emerald is grue.
>
> The question is whether we should conjecture that all emeralds are green rather than that all emeralds are grue when we obtain a sample of green emeralds examined before time *t*, and if so, why.

Clearly we have a family of grue predicates in this problem, corresponding to different "critical times" *t*; let's write grue(*t*) to denote these. Following Goodman, let us refer to methods as projection rules in discussing this example. A projection rule succeeds in a world just in case it settles on a generalization that is correct in that world. Thus in a world in which all examined emeralds are found to be green, we want our projection rule to converge to the proposition that all emeralds are green. If all examined emeralds are grue(t), we want our projection rule to converge to the proposition that all emeralds are grue(t). Note that this stipulation treats green and grue predicates completely on a par, with no bias towards either. As before, let us consider two rules: the "natural" projection rule which conjectures that all emeralds are green as long as only green emeralds are found, and the "gruesome" rule which keeps projecting the next grue predicate consistent with the available evidence. Expressed in the green-blue vocabulary, the gruesome projection rule conjectures that after observing some number of n green emeralds, all future ones will be blue. The figure below illustrates the possible observation sequences and the natural projection rule in this model of the New Riddle of Induction.

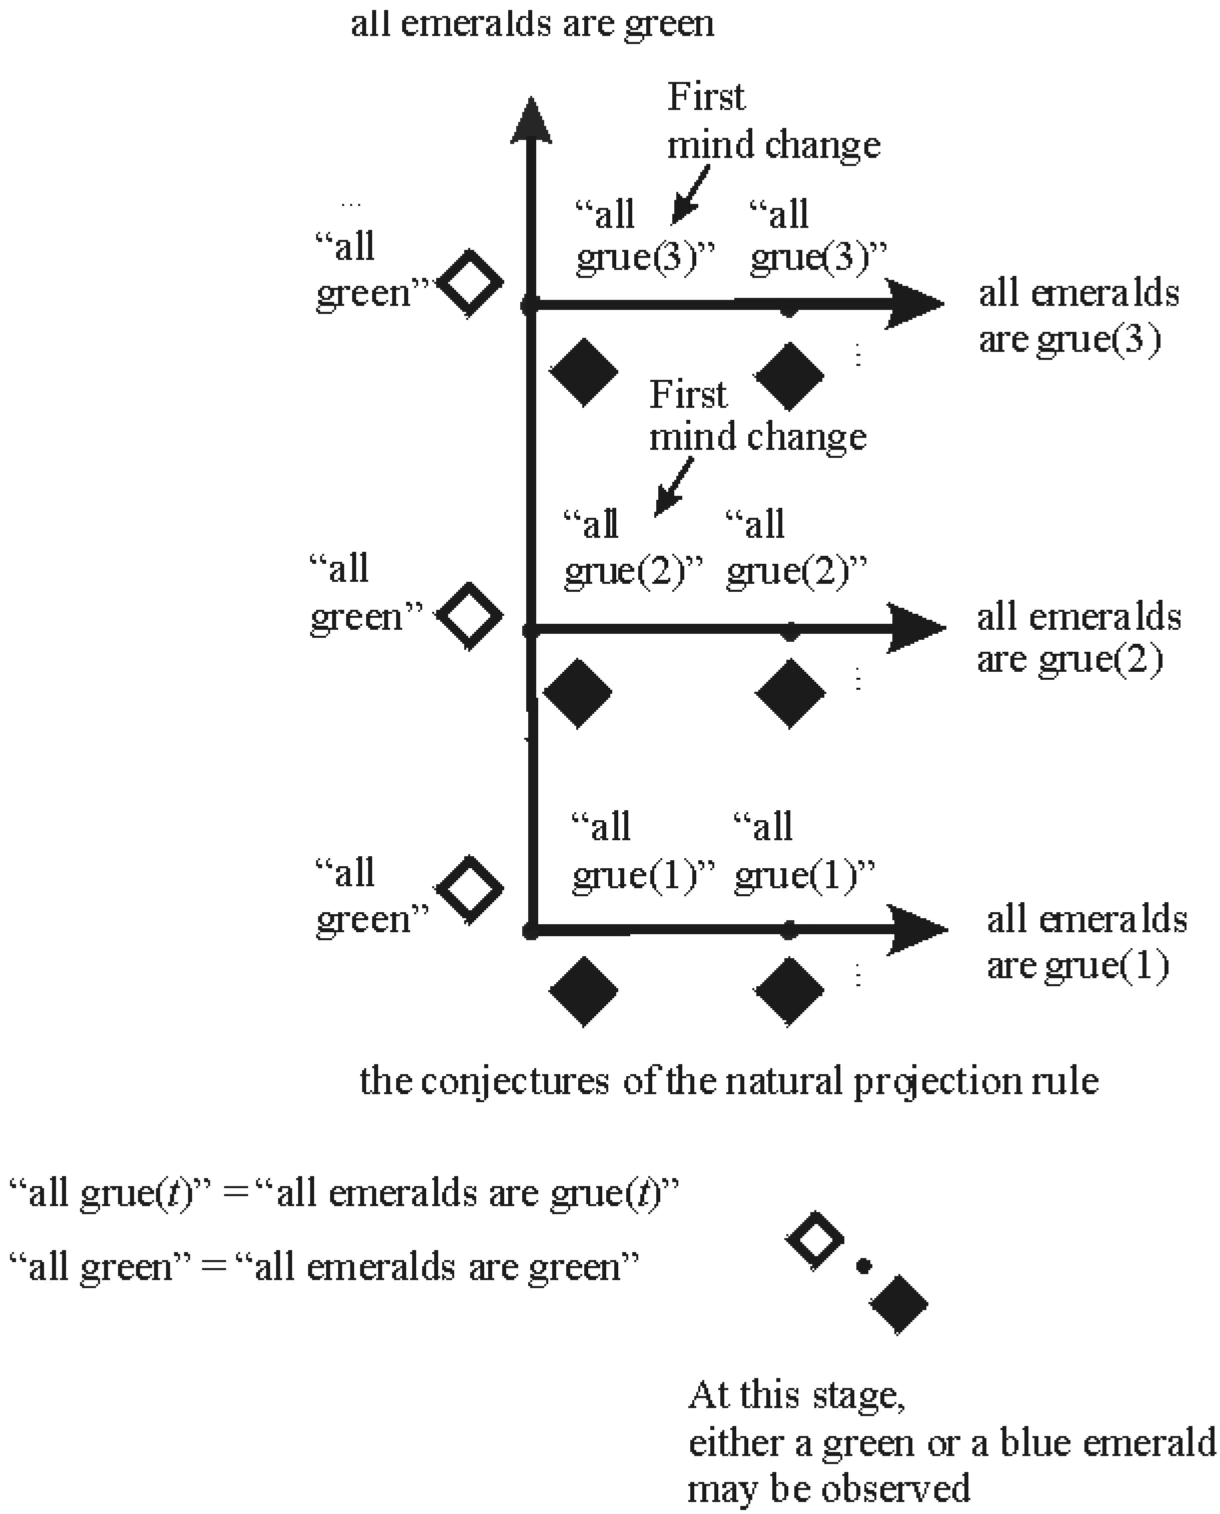

The following figure shows the gruesome projection rule.

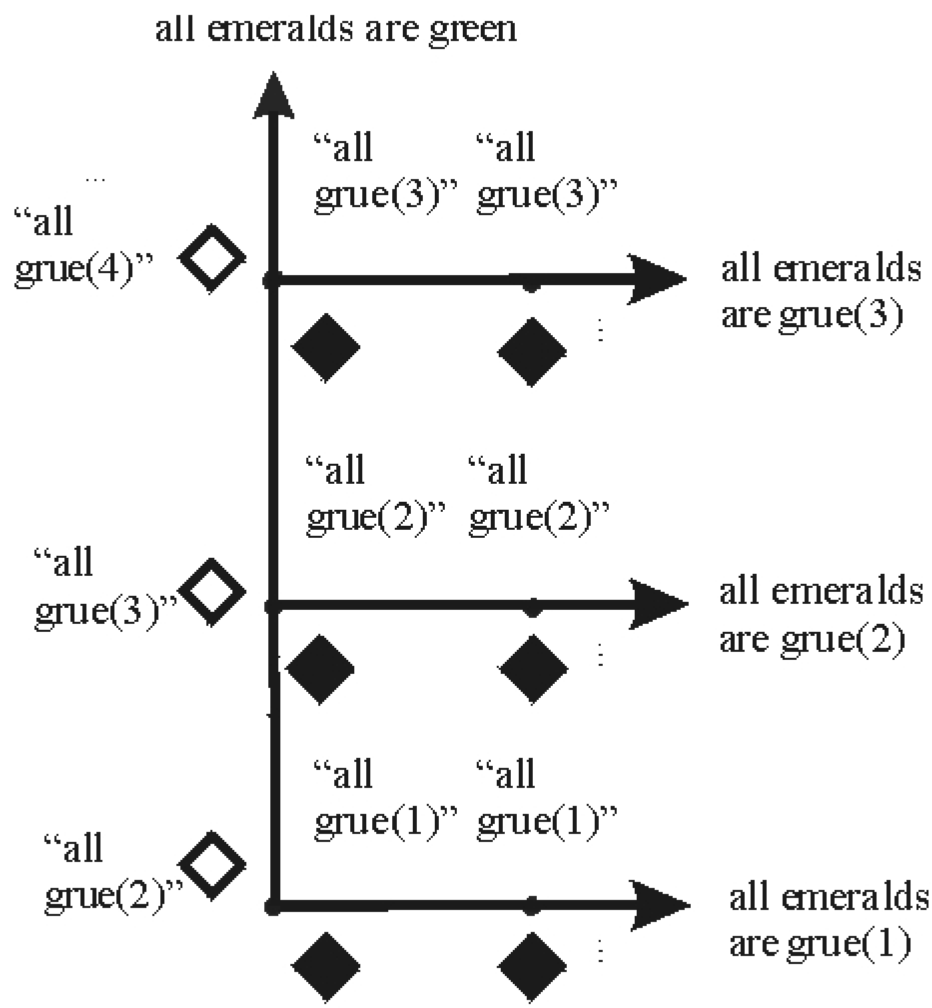

the conjectures of the gruesome projection rule

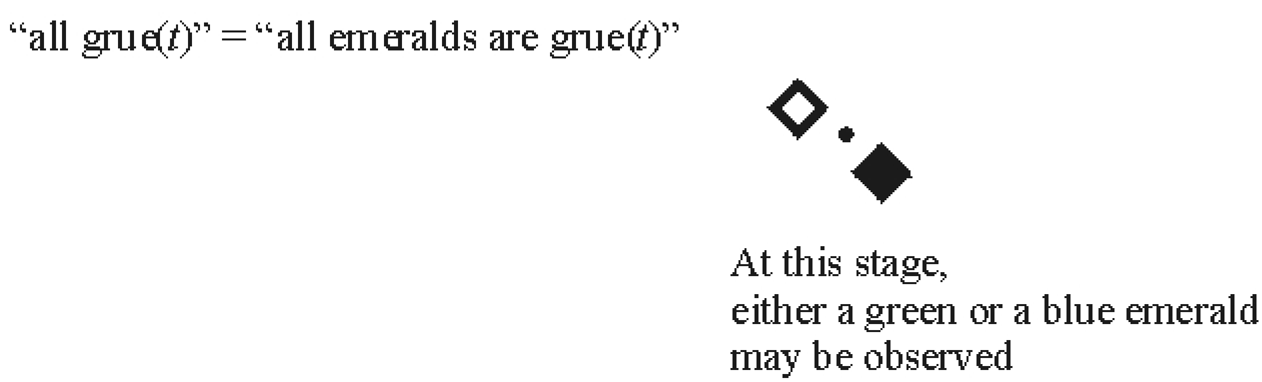

How do these rules measure up to the goal of arriving at a true generalization? Suppose for the sake of the example that the only serious possibilities under consideration are that either all emeralds are green or that all emeralds are grue(*t*) for some critical time *t*. Then the natural projection rule settles on the correct generalization no matter what the correct generalization is. For if all emeralds are green, the natural projection rule asserts this fact from the beginning. And suppose that all emeralds are grue(*t*) for some critical time *t*. Then at time *t*, a blue emerald will be observed. At this point the natural projection rule settles on the conjecture that all emeralds are grue(*t*), which must be correct given our assumption about the possible observation sequences. Thus no matter what evidence is obtained in the course of inquiry -- consistent with our background assumptions -- the natural projection rule eventually settles on a correct generalization about the colour of emeralds.

The gruesome rule does not do as well. For if all emeralds are green, the rule will never conjecture this fact because it keeps projecting grue predicates. Hence there is a possible observation sequence -- namely those on which all emeralds are green -- on which the gruesome rule fails to converge to the right generalization. So means-ends analysis would recommend the natural projection rule over the gruesome rule. Some comments are in order.

1. As in the previous example, nothing in this argument hinges on arguments to the effect that certain possibilities are not to be taken seriously a priori. In particular, nothing in the argument says that generalizations with grue predicates are ill-formed, unlawlike, or in some other way a priori inferior to "all emeralds are green".

2. The analysis does not depend on the vocabulary in which the evidence and generalizations are framed. For ease of exposition, I have mostly used the green-blue reference frame. However, grue-bleen speakers would agree that the aim of reliably settling on a correct generalization requires the natural projection rule rather than the gruesome one, even if they would want to express the conjectures of the natural rule in their grue-bleen language rather than the blue-green language that we have used so far. (For more on the language-invariance of means-ends analysis see Section 4 (The Limits of Inquiry and the Complexity of Empirical Problems), as well as Schulte [1999a, 1999b]).

3. Though the analysis does not depend on language, it does depend on assumptions about what the possible observation sequences are. The example as described above seems to comprise the possibilities that correspond to the colour predicates Goodman himself discussed. But means-ends analysis applies just as much to other sets of possible predicates. Schulte [1999a, 1999b] and Chart [2000] discuss a number of other versions of the Riddle of Induction, in some of which means-ends analysis favours projecting that all emeralds are grue on a sample of all green emeralds.

4. Even with the assumptions granted so far, there are reliable projection rules that project that all emeralds are grue(*t*) on a sample of all green emeralds. For example, the projection rule "conjecture that all emeralds are grue(3) until 3 green emeralds are observed; then conjecture that all emeralds are green until a blue emerald is observed" is guaranteed to eventually settle on a correct generalization just like the natural projection rule. (It's a worthwhile exercise to verify the reliability of this rule.) I will discuss criteria for further restricting the space of rules in Section 3 (The Long Run in The Short Run).