

ENTRENCHMENT AND THE KRIPKE-WITTGENSTEIN PARADOX

Chris Dorst

A thesis submitted to the faculty at the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Master of Philosophy in the Department of Philosophy.

Chapel Hill
2014

Approved by:

Marc Lange

Matt Kotzen

John Roberts

© 2014
Chris Dorst
ALL RIGHTS RESERVED

ABSTRACT

Chris Dorst: Entrenchment and the Kripke-Wittgenstein Paradox
(Under the direction of Marc Lange)

Surprisingly little attention has been paid to Goodman's own entrenchment-based solution to his New Riddle of Induction. Ostensibly, this is because Goodman views the problem presented by the New Riddle differently than most contemporary philosophers; Goodman sees the problem as one of codifying the acceptable inductive inferences rather than justifying why those inferences are acceptable. I argue that if we share his perspective, Goodman's own solution to the New Riddle is undermined: the linguistic facts about the entrenchment of predicates are no more accessible than facts about which classes are relevant to nature. The result is that Goodman's arguments leave us in a predicament very much like that presented by the Kripke-Wittgenstein paradox about meaning.

ACKNOWLEDGEMENTS

I owe a great deal of thanks to many people, including Marc Lange, Matt Kotzen, John Roberts, Douglas Stalker, Kayla Bishop, and Kevin Dorst for help in sharpening my thoughts about the New Riddle, reading drafts, and providing feedback at every stage of the process.

Thanks also to Frederick Eberhardt, who first introduced me to the New Riddle and impressed its seriousness upon me; to Joseph Ullian and Catherine Elgin, who patiently answered many questions; to my dad, Stanley Dorst, with whom I had many discussions that shaped my view of the New Riddle; and to my mom, Margaret Yoest, for putting up with those discussions.

TABLE OF CONTENTS

Section 1: Introduction.....	1
Section 2: Goodman's Problem.....	3
Section 3: Kripke-Wittgenstein's Problem.....	8
Section 4: Goodman's Solution.....	12
Section 5: Problematizing Entrenchment.....	29
Section 6: Conclusion.....	42
REFERENCES.....	46

1 Introduction

Considering all of the weight that gets attached to Nelson Goodman's New Riddle of Induction, it is curious that, of all the proposed solutions, it is Goodman's own solution that gets treated perhaps the least seriously. Mere entrenchment, or "history of use," is often thought to be too accidental, too much a matter of chance, to serve as the foundation for an inductive logic.¹ It is rare, though, that arguments of any appreciable rigor are advanced against it.² Perhaps this is because it is taken to be just obviously inadequate; Goodman's theory does not even have the appropriate *form* of a solution. Without getting too far into the theory itself (see §4), the common thought is roughly that, at best, Goodman provides us with a *description*, where what we wanted was a *justification*. At any rate, that is the impression that one often gets.

Part of my purpose in this paper is to suggest that Goodman's entrenchment solution to his New Riddle deserves more serious consideration than we have, in general, been inclined to give it. That we have tended to undervalue his solution is no surprise; Goodman's way of thinking is easily misunderstood, and there are arguments lurking just beneath the surface of his writings that have to be pieced together in order to fully appreciate his motivations. I realize this all sounds rather cryptic at the moment, but it should become clearer in §2, and especially §4,

¹There are notable exceptions: Robert Schwartz and Israel Scheffler (1970), Joseph Ullian (personal communication), Catherine Elgin (1988), and Ian Hacking (1994) all express sympathies with the entrenchment solution.

²Again, there are exceptions: see Zabłudowski's series of ingenious attacks on Goodman's theory, starting with his (1974) and extending through his (1982). Zabłudowski's charges are met with equally brilliant rebuttals in Goodman and Ullian's series of replies. For a summary of their exchanges, see the immensely useful annotated bibliography in Stalker (1994).

when I outline the problem and his solution in more detail.

Lending credence to Goodman's entrenchment solution, however, is only a part of my purpose here. My broader goal is to elucidate a deep connection between Goodman's view of his own problem and the central skeptical paradox of Kripke's *Wittgenstein on Rules and Private Language*. My claim is that if we share Goodman's motivations with respect to the New Riddle, then two surprising results present themselves: (1) despite the credence I intend to lend it, Goodman's entrenchment solution itself is undermined by those very motivations, and (2) the New Riddle can be seen as a special case of the paradox presented by Kripke-Wittgenstein (sometimes 'KW'³). There is of course a serious question of whether we ought to share Goodman's motivations with respect to the New Riddle, and while I will suggest that there are some surprising ways to substantiate his view, ultimately I think it is an open question whether he is right.

In the next section, I outline the problem(s) that Goodman has posed for us with the New Riddle, and in §3 I outline KW's "skeptical paradox" about meaning. In §4 I present Goodman's solution to the New Riddle, paying special attention to his own, somewhat veiled, motivations for the solution. In §5, I argue that Goodman's own motivations, though they cast his entrenchment solution in a new light, ultimately undermine that very solution. They do so in such a way that we are left, if we accept them, right in the middle of KW's paradox. Thus, as I will suggest later, there is a strong sense in which Goodmanian considerations undermine not only the entrenchment solution, but the very posing of the New Riddle as well. In §6 I conclude

³I follow the custom of attributing the paradox to neither Kripke nor Wittgenstein directly, but to the reading of Wittgenstein "as it suggested a problem" to Kripke. That being said, I am inclined to think that the paradox is a sharp way of posing something close to Wittgenstein's motivations in *Philosophical Investigations*. But that is a discussion for another time.

by considering where this leaves us: ought we accept Goodman's arguments, or can the skeptical conclusions be avoided?

2 Goodman's Problem

Goodman first presented his famous problem in his (1946), but its most familiar presentation came nine years later in *Fact, Fiction, and Forecast* (pp. 72-83).⁴ In reviewing the problem and its surrounding dialectic, I will be treading over very familiar ground for most readers. I have to plead the reader's patience; the pedantry here will pay off when I attempt to clarify Goodman's motivations later on.

The problem, then, is as follows. Most people would accept as valid⁵ the following inductive inference (*pace* worries about the old problem of induction):

1. All observed emeralds are green at the time they are observed.

∴ All emeralds are always green.

Now suppose we define a new predicate, 'grue', as follows: something is grue at a time if it is green at that time and the time is before midnight on January 1st of the year 2020, or it is blue at that time and the time is after midnight on January 1st, 2020.⁶ Most people would *not* accept the following inference as valid:

2. All observed emeralds are grue at the time they are observed.

∴ All emeralds are always grue.

⁴Citations are from the Fourth Edition, published in 1983.

⁵I am adopting Goodman's somewhat idiosyncratic usage of "valid," not the notion that we traditionally associate with deductive inference. He sometimes describes himself as looking for a criterion of inductive validity, e.g. on p. 84 of *Fact, Fiction, and Forecast*.

⁶Here I am using a slightly modified definition, which has the virtue of simplifying matters compared to Goodman's original definition by avoiding reference to observation.

Inference (2) would lead us to think that any emeralds observed after 2020 would be grue, and therefore blue, at that time, whereas inference (1) would lead us to think that emeralds observed after 2020 would be green at that time. The problem, of course, is to say what it is about inference (1) that makes it acceptable and what it is about inference (2) that makes it unacceptable. Given that the only apparent difference between the inferences lies in the predicates used in each, it is natural to think that our attention should be focused there, and indeed that is where Goodman and many others have looked for solutions. That is, many attempted solutions to the New Riddle try to show why inferences like (1) are acceptable and inferences like (2) are unacceptable by showing that there is something defective about predicates like ‘grue’ (or something superior about predicates like ‘green’).

In that vein, Goodman considers several proposals as to why inferences like (1) are valid and inferences like (2) are not. Most germane to our purposes is his treatment of the suggestion that predicates like ‘green’ are ‘nonpositional’ or ‘qualitative’, whereas this is not true for predicates like ‘grue.’ Intuitively, if a predicate is positional, i.e. it involves reference to a particular time, then it is not suited for use in induction across that particular time. To this Goodman replies, first, that it is not at all clear to him how to tell whether a predicate is ‘nonpositional’ or ‘qualitative’ except by completely begging the question and asking whether the predicate is inductively ‘well-behaved’. In other words, the only criterion we have by which to judge whether a predicate is qualitative, for example, is to figure out whether it figures in inductions that we find acceptable. But of course that is exactly the notion we are trying to get a grip on, i.e. which inductions are acceptable.

Whether or not one agrees that our grip on predicate positionality is dependent on our grip on acceptable inductions, Goodman’s second reply to this suggestion is ostensibly far more

damning. To the claim that the definition of ‘grue’ requires reference to time (demonstrating that it is a ‘positional’ predicate), he responds by considering another new predicate, ‘bleen.’ Something is bleen at a time if it is blue at that time and the time is before midnight on January 1st, 2020, or it is green at that time and the time is after midnight on January 1st, 2020.⁷ Now it is true that we can define ‘grue’ in terms of green and blue and a reference to time, but it is equally true that we can define ‘green’ in terms of grue and bleen and a reference to time: something is green at a time if it is grue at that time and the time is before 2020 or it is bleen at that time and the time is after 2020. So whether ‘grue’ or ‘green’ counts as positional or qualitative depends on what language we start from. As Goodman puts it, “qualitativeness is an entirely relative matter and does not by itself establish any dichotomy of predicates” (p. 80).

There are of course a variety of other proposed solutions to the New Riddle, and I will consider several more in this paper. For now, the preceding paragraphs should serve as a brief reminder of the problem and some of the dialectic surrounding it. It is important to note, however, that referring to it as a ‘problem’ (singular) is in fact misleading. There are two distinct problems here. First there is the task of *demarcating* the acceptable (‘projectible’) inductive inferences from the unacceptable (‘unprojectible’) ones. As I mentioned above, proposals to this effect often focus on demarcating the projectible *predicates*. Second there is the task of *justifying* why the acceptable inferences are acceptable. What is it about the demarcating criterion that makes exactly those hypotheses (or predicates) that possess it the projectible ones?

It is easy to run these two tasks together, for we are naturally inclined to be more concerned with the second than the first, and the task of justification presupposes that we have already demarcated the hypotheses whose suitability for induction we are now attempting to

⁷I will abbreviate this date as ‘2020’ henceforth.

elucidate. But one can demarcate without justifying, and Goodman himself seems to be more concerned with the former than the latter. Consider what he says at the beginning of the next section (suggestively titled “A New Look at the Problem”):

The problem of confirmation, or of valid projection, is the problem of *defining* a certain relationship between evidence or base cases on the one hand, and hypotheses, predictions or projections on the other...What we want is an accurate and general way of *saying which* hypotheses are confirmed by, or which projections are validly made from, any given evidence. (1983, p. 84, my emphasis)

My emphasis in the above is meant to point out that his concern is primarily descriptive, not normative. Our interest, he says, is with *definition*, with *saying which* hypotheses are confirmed by given evidence,⁸ not *saying why* those hypotheses are so confirmed. Ullian (1961a) likewise stresses that Goodman is fundamentally looking for “codification” of the acceptable inferences, not justification.

There is admittedly a question of whether Goodman would assent to this way of stating his concern. His conception of epistemic justification is notoriously lightweight:

How do we justify a *deduction*? Plainly, by showing that it conforms to the general rules of deductive inference...Rules and particular inferences alike are justified by being brought into agreement with each other. *A rule is amended if it yields an inference we are unwilling to accept; an inference is rejected if it violates a rule we are unwilling to amend...*and in the agreement achieved lies the only justification needed for either. All this applies equally well to induction...Predictions are justified if they conform to valid canons of induction; and the canons are valid if they accurately codify accepted inductive practice. (1983, pp. 63-64, emphasis in original)

For Goodman, then, what makes the rules of induction (whatever they may be) valid is just that

⁸The notion of ‘given evidence’ is a bit of a loose way of speaking, according to Goodman. As he puts it, “evidence consists of statements, and the statements ‘emerald a is green’, etc. that confirm the [green] hypothesis are not the same as the statements ‘emerald a is grue’, etc. that confirm the [grue hypothesis]” (1972, p. 359). So it is not strictly speaking correct to say that the New Riddle requires us to say which projections are validly made from given evidence, unless we take the “given evidence” to consist of something like an equivalence class of evidence statements. Perhaps it would be more accurate to say that the New Riddle requires us to say which evidence statements are to be privileged.

they conform to accepted inductive practice. So by articulating a rule that describes the acceptable inductions, we are simultaneously demonstrating the validity of that rule. This is just to emphasize that Goodman would probably not agree with the sharp distinction I have drawn between demarcation and justification. Justification is demarcation, according to Goodman. To ask for a more robust kind of justification is to ask a pseudo-question. Of course, I take it that most people do not share this view of epistemic justification, which is why I have emphasized the distinction. That is, most philosophers would expect more from an answer as to why *these* particular inductive inferences are valid than the rather brusque reply that they accord with such-and-such a rule, and that rule accurately codifies the inductive inferences we take to be valid.

Now, there is a sense in which this cannot be the whole story here. Goodman's predicate 'projectible' picks out exactly the hypotheses we take to be inductively confirmable (when construed as a predicate of hypotheses) or exactly the predicates we take to figure in valid inductive inferences (when construed as a predicate of predicates). Presumably, it would not solve the New Riddle just to say that the valid inductions are the ones that use only projectible predicates. But why not? After all, that rule itself exactly demarcates the valid inductions. Well, is the rule of any use to us? We check whether an induction is valid by checking whether it contains only projectible predicates, and what do we do to check whether its predicates are projectible? Perhaps the problem with using 'projectible' as a solution to Goodman's Riddle is analogous to the first problem he mentions with using 'qualitative': our grip on the projectible predicates is *via* our grip on which predicates figure in valid inductions. Descriptions can be circular just like justifications. What we want (or at least what Goodman wants) is an *informative* description of the valid inductions, and the flat-footed reply that "the valid inductions are the ones that use only projectible predicates" is anything but informative. Compare Goodman's

description of the unhappy state of confirmation theory given the New Riddle: "...what we have in current confirmation theory is a definition that is adequate for certain cases that so far can be described only as those for which it is adequate. The theory works where it works. A hypothesis is confirmed by statements related to it in the prescribed way provided it is so confirmed" (pp. 81-82). Clearly, Goodman does not think that just any demarcation will do.

Enough Goodman exegesis (for now). The important consideration to note at this point is that there are two problems that go by the name of the "New Riddle." Ask a Goodmanian what the problem is and he will say that it is to demarcate the acceptable inductions from the unacceptable ones. Ask your average philosopher on the street and they will likely construe it as primarily a justificatory one, where "justification" involves more than just informative codification. It is crucial to keep this distinction in mind when considering whether a proposed solution actually works, and when considering why Goodman himself rejects certain solutions. Since the task of justification (as we ordinarily conceive it) involves strictly more than the task of demarcation, the fact that a proposed solution fails at the former does not immediately imply that it also fails at the latter.

3 Kripke-Wittgenstein's Problem

Leave behind Goodman, for a moment; I want to briefly review the central skeptical paradox of Kripke's *Wittgenstein on Rules and Private Language* (1982). This interlude may appear irrelevant at present, but it will become very relevant shortly.

Like Goodman, Kripke motivates his problem by focusing on a particular example.⁹ Suppose I have never computed the sum of any numbers greater than or equal to 57, and now I am asked to compute the sum of 68 and 57. I do so, and I come up with the answer '125'. Now a

⁹The example is first introduced on pp. 8-9.

“bizarre skeptic” comes along and questions whether I am using the word ‘plus’ (or ‘sum’) with the same meaning that I did in the past. In fact, he suggests that in the past when I used the word ‘plus’, I really meant the “quus” function, which is denoted by ‘ \oplus ’ and defined as follows:

$$\begin{aligned} x \oplus y &= x + y && \text{if } x, y < 57 \\ &= 5 && \text{otherwise} \end{aligned}$$

Thus the skeptic thinks that if I am to accord with my previous meaning, I should answer ‘5’ rather than ‘125’ when I am prompted for the sum of 68 and 57. According to Kripke, the challenge is now to point to some fact about my past usage of the term ‘plus’ that makes it the case that I really meant plus and not quus. By hypothesis, I never actually computed the sum of any numbers greater than 57, so I cannot point to any such occasion to settle the matter. So what fact *can* I point to?

The tempting response is just to say that I meant plus rather than quus because of certain instructions I gave myself regarding addition. What these instructions amount to, though, is very difficult to articulate. The instructions could not be said to include the explicit directive to “answer ‘125’ if prompted to compute 68+57,” because we have supposed that I never computed the sum, nor have I explicitly thought about such a computation. (If I *have* thought about it, we can always change the example to one that I have not considered.) And it is likewise of no help to suggest that my instructions consisted in giving myself a few examples of addition and telling myself to “go on in the same way” for the computation of other sums; *both* addition and quaddition could be construed as “going on in the same way,” as they are equally consistent with all of the examples I had previously given myself.

A more promising response would be to say that the instructions consisted in some *rule* that determines how addition is to be continued beyond the examples that I have computed. Such

a rule would have to be markedly more detailed than the vague instruction to “go on in the same way,” which essentially amounts to nothing more than “and so on.” Kripke considers several proposals for such a rule, but it is sufficient here to just consider the first, which is as follows. Suppose we are adding x and y . To do so, count out x marbles and put them in a heap. Then count out y marbles and put them in another heap. Then put these two heaps together and count the marbles in the joint heap. This result is the sum of x and y .

Call this rule the ‘Marble Rule.’ Obviously we do not actually use the Marble Rule to compute sums, but its excessive simplicity is instructive in this case. So suppose this was the rule I articulated to myself. Would I not then have definitively determined that I meant plus and not quus? Here is Kripke’s response on behalf of the skeptic:

True, if ‘count’, as I used the word in the past, referred to the act of counting (and my other past words are correctly interpreted in the standard way), then ‘plus’ must have stood for addition [as opposed to quaddition]. But I applied ‘count’, like ‘plus’ to only finitely many past cases. Thus the skeptic can question my present interpretation of my past usage of ‘count’ as he did with ‘plus’. In particular, he can claim that by ‘count’ I formerly meant *quount*, where to ‘quount’ a heap is to count it in the ordinary sense, unless the heap was formed as the union of two heaps, one of which has 57 or more items, in which case one must automatically give the answer ‘5’. (p. 16)

If we interpret ‘count’ in this nonstandard way, and we suppose that the Marble Rule determined my meaning of ‘plus’, the result will be that I meant quaddition rather than addition. So the appeal to the Marble Rule has really only pushed our problem back a step; now we are left wondering what fact constitutes my using ‘count’ to mean count rather than quount. We may then appeal to another rule to fix the meaning of ‘count,’ but by now it should be clear that we can go on in the same way with respect to this new rule as well. That is, given a rule that fixes the meaning of a particular word, we can always give a nonstandard interpretation to some word(s) in the articulation of the rule itself and ask what makes it the case that we mean the standard meaning as opposed to the nonstandard one. And the same considerations would

obviously apply if we used some rule other than the Marble Rule in an attempt to fix the meaning of ‘addition.’ (As I mentioned, the simplicity of the Marble Rule pays off here. If a rule *that* simple can be problematized by reinterpretation, then surely more complicated rules can as well.)

Kripke goes on to consider several other proposals for facts that could constitute my meaning plus rather than quus, and he rejects them all for varying reasons. Fortunately we need not rehearse the entire dialectic here. What is important is the upshot: it begins to look like there is no fact that could make it the case that I meant plus and not quus with my past uses of ‘plus.’ Indeed, there appears to be no fact that could distinguish between my meaning plus and my meaning nothing at all in the past, for if I meant *anything*, surely I meant plus rather than quus. If so, Kripke is quick to point out that there is likewise no such fact about my *present* use of ‘plus.’ And now, as Kripke puts it, “it seems that the entire idea of meaning vanishes into thin air” (p. 22).

So much for our brief review of KW’s skeptical paradox. Ostensibly, the paradox has little to do with Goodman’s New Riddle. Granted, both puzzles involve unnatural neologisms (‘grue’ and ‘quus’) that cause trouble for us, but that is about where the (superficial) similarities stop. The problems posed by these neologisms are quite different: on the one hand we have a problem about prediction (which regularities should we project into the future?), and on the other hand we have a problem about meaning (what makes it the case that I mean *x* as opposed to *y*, or as opposed to *nothing*, by my use of a word?). But there is at least this further similarity: in both cases, we naïvely expect some instruction like “go on in the same way” to get the job done. To the question “Which regularities should we project into the future?” we expect to be able to respond: “Look at the regularities that have held in the past, and project those same regularities into the future.” And to the question “Which regularities in past uses of ‘plus’ should I extend

into my future uses of ‘plus’ (so as to use ‘plus’ with the same meaning as I did in the past)?” we expect to be able to answer: “You followed Rule *X* in your application of ‘plus’ before; if you go on following Rule *X*, you will mean the same thing.” In both cases, it is some nonstandard

4 Goodman’s Solution

I have stressed that Goodman conceives of the New Riddle as a problem of demarcation: which predicates are acceptable for use in induction?¹⁰ I have also stressed that it is natural to view the problem as a justificatory one, where we mean something more robust than mere demarcation by ‘justification.’ With this in mind, we ought to be careful when interpreting Goodman’s rejections of proposed solutions to his Riddle. Consider, for instance, his (second) rejection of the ‘positionality’ criterion discussed above. It was natural to read that as follows: we cannot appeal to nonpositionality to justify why the projectible predicates are projectible, because grue-speakers would have the same justification open to them. That is, if we start from *our* language then it is clear that ‘grue’ and ‘bleen’ are positional, but if we start from the language of the grue-speakers, then it is clear that ‘green’ and ‘blue’ are positional. The symmetry of the situation thus renders our attempt at justification ineffectual; grue-speakers would be able to give the same justification for the projectibility of *their* predicates that we are attempting to give for ours.

Since Goodman is concerned only with demarcation, this cannot be how he views the problem with an appeal to positionality. If we are looking merely to mark off a class of predicates, then what could be problematic about saying that the projectible predicates are the

¹⁰Framing the question as one of demarcation suggests that what Goodman is really asking for is something like a conceptual analysis of our concept of ‘projectible hypothesis.’ Given the apparent lack of any such analysis for most concepts (‘game,’ ‘causation,’ ‘knowledge,’ ‘chair,’ etc.), we should perhaps feel *lucky* if there is one for ‘projectible hypothesis,’ rather than unlucky if there is not.

ones that we have to define with reference to time? One might think that we have an intuitive enough grip on whether a definition of a predicate in our language would necessarily involve reference to time, and if so, then it is unclear why positionality would fail even as a demarcation criterion. The fact that grue-speakers would use it to mark off a different class of predicates is plainly *irrelevant* to whether *we* can demarcate the projectible predicates by appeal to positionality.¹¹

To get clearer on why Goodman thinks the appeal to positionality actually fails, it will help to have a view of his solution.¹² Goodman proposes to demarcate the projectible predicates by appeal to entrenchment: the projectible predicates are the ones that are most entrenched, i.e. have the longest histories of projective use in our linguistic community.¹³ The projectible hypotheses are then (roughly) the ones that contain only entrenched predicates. More specifically, consider only the universal conditional hypotheses that are supported (i.e. have positive instances), unviolated (i.e. have no negative instances), and unexhausted (i.e. are not yet entailed

¹¹I am here ignoring the question of whether all and only the projectible predicates would be nonpositional, but that proposition seems dubious. Consider Goodman's own example: if we sample a few of the men in a given room and find that they are all third sons, this does not confirm (to any appreciable degree) the hypothesis that all men in that room are third sons. Here the predicate 'is a third son' does not involve any obvious reference to time in the problematic way that 'grue' does. Of course 'positionality' might be construed more broadly, encompassing reference to a particular position not just in time but in space, birth order, etc. This would avoid the worry that there might be unprojectible predicates that are nonpositional, but it would incur the opposite worry that there might be projectible predicates that *are* positional. I ignore these considerations here because Goodman himself ignores them in his treatment of the positionality proposal.

¹²See pp. 84-124 of *Fact, Fiction, and Forecast* for Goodman's explication of the theory.

¹³Goodman admits that his distinction between the projective use and the non-projective use of a predicate is inexact (1983, pp. 88-89). For this reason, I am inclined not to read too much into the requirement that we count only projective uses; if anything important hinges on it, Goodman should have been more careful here.

by the evidence). The green and grue hypotheses about emeralds would both be among this set. Entrenchment is used to decide among hypotheses in this set that conflict: the projectible hypotheses are the ones with significantly better entrenched predicates in either the antecedent or consequent position, and at least equally-well entrenched predicates in the opposite position. So in comparing (1) ‘All emeralds are always green’ with the conflicting hypothesis (2) ‘All emeralds are always grue,’ we note that the antecedent predicate, ‘emerald,’ is the same in each case, whereas the consequent predicate of (1) is significantly more entrenched than the consequent predicate of (2). Thus (1) is the projectible hypothesis, and (2) is deemed unprojectible. So entrenchment enters, as it were, as a tiebreaker between conflicting hypotheses that are supported, unviolated, and unexhausted.

It is critical to note that Goodman actually regards entrenchment as a property of *classes*, i.e. predicate extensions, rather than of predicates themselves:

The entrenchment of a predicate results from the actual projection not merely of that predicate alone but also of all predicates coextensive with it. In a sense, not the word itself but the class it selects is what becomes entrenched, and so to speak of the entrenchment of a predicate is to speak elliptically of the entrenchment of the extension of that predicate. (1983, p. 95)

That being said, entrenchment is supposed to be an entirely unproblematic criterion, epistemically speaking. Learning facts about entrenchment is straightforward, according to Goodman. We measure the entrenchment of a class by counting the past history of projections using predicates that denote that class. Clearly, the class of green things has a much longer history of projective uses—it has been denoted in many more actually projected hypotheses—than the class of grue things. This is not to say that any speaker of a given language automatically knows all of the facts about entrenchment, only that it is typically unproblematic to learn such facts.

With Goodman's solution on the table, let us return to the question of why he rejects the positionality proposal. It will be instructive here to examine his response to Barker and Achinstein (1960), who attempt to get more mileage out of the positionality criterion as a solution to the New Riddle. We need not review their entire paper here, but their proposal essentially comes down to the claim that the projectible predicates are the ones that can be represented with a single picture. For example, one could represent green by using paint of only one pigment. Conversely, they maintain that if someone, even a grue- speaker, were attempting to represent grue, she would have to use paint of two different pigments: green pigment would represent grueness before 2020 and blue pigment would represent grueness afterward. And it is clear that they take nonpositionality not just to demarcate, but to justify why the projectible predicates are projectible.¹⁴

Goodman (1960) replies by stressing the flexibility of our notion of representation. We could, for example, use black and white cross-hatching to represent a grue object, and we could do so in a single picture. Of course, this would be a somewhat nonstandard mode of representation, but to object to it on these grounds would be a mistake, as Goodman goes on to point out:

[If Barker and Achinstein object to nonstandard modes of representation] they are saying, then, that so long as only accustomed modes of representation...are used, a non-positional predicate is one such that there is a single time-indifferent representation for all its instances of application. But if one is willing to put the matter thus, relying upon restriction to a familiar representational vocabulary, representation need not be brought in at all. "Grue" can be classed as positional on the grounds that in terms of ordinary, familiar language, no one term describes the color of all grue objects; two different ordinary words, "green" and "blue," are needed. This, however, depends entirely upon "green" and "blue" rather than "grue" happening to belong to ordinary language; that is,

¹⁴"[If we are correct]...a logically important difference does exist between 'grue' and 'green,' a difference sufficiently important to help us see why it is that we rightly regard 'green' as more projectible than 'grue.'" (p. 521)

it depends entirely upon the facts of habit or entrenchment. In resting their definition of positionality upon a restriction to the most accustomed means of representation, Barker and Achinstein are making much the same appeal to entrenchment but in a more roundabout and covert way. (p. 524)

I want to carefully consider this response, because I think it contains a deep-seated motivation for Goodman's entrenchment solution. If we try to use 'representable with a single picture' to pick out the projectible predicates, Goodman will press us to consider nonstandard modes of representation. True, green is representable with a single picture and grue is not *according to standard modes of representation*. But if we allow ourselves nonstandard modes, then we have not demarcated the class of predicates that was our goal—*any* predicate can be represented by a single picture in that case. So for their criterion to demarcate exactly the predicates it is meant to, Barker and Achinstein must restrict the modes of representation under consideration to the traditional ones. This, admittedly, might serve as an acceptable demarcation of the projectible predicates. The problem is that we have essentially appealed to entrenchment in stating it; in restricting our attention to traditional modes of representation, we have made ineliminable reference to our traditional use of 'representation.' If such reference is acceptable, we might as well have just appealed to the fact that we traditionally denote the classes of green things and blue things rather than the class of grue things, i.e. to the fact that 'green' and 'blue' are entrenched in our language whereas 'grue' is not. Appealing to traditional modes of representation is thus merely a circuitous way of appealing to tradition of use, i.e. entrenchment. In other words, if we are appealing to tradition of our use of 'representation,' then we might as well just appeal to tradition of use in general to distinguish the projectible predicates from the nonprojectible ones.¹⁵

¹⁵I should mention that this passage indicates that Goodman himself drops the requirement that only projective uses count toward the entrenchment of a predicate. Here he essentially equates

In summary, the dialectic is as follows. Barker and Achinstein propose a criterion, nonpositionality, that they think both demarcates the projectible predicates and justifies their projectibility. Goodman responds that nonpositionality does not even work as a demarcation criterion unless we understand it in the traditional way, which is essentially just to appeal to entrenchment. This, I expect, is ultimately the reason that Goodman thinks an appeal to qualitiveness fails. That “qualitiveness is an entirely relative matter” does not seem to be immediately problematic if we are only trying to *demarcate* by appeal to it. But the relativity is problematic because it requires us to understand it in the traditional manner, and thus the appeal to qualitiveness boils down to a circuitous appeal to tradition of use.

Independent of the fact that Goodman gives this reply to Barker and Achinstein, there is good reason to think that these considerations are central to his way of viewing the New Riddle in the first place. I say this even while noting that they do not explicitly appear in *Fact, Fiction, and Forecast*. One might object that since Goodman motivates the entrenchment solution independently of these considerations in *Fact, Fiction, and Forecast*, his reply to Barker and Achinstein can be nothing more than an ancillary vindication of that solution. I sympathize with the spirit of this objection, but I suspect that it contains a false presupposition. In other words, I do not think that Goodman motivates the entrenchment solution independently of the considerations presented in his reply to Barker and Achinstein. If we carefully attend to his arguments in *Fact, Fiction, and Forecast*, we are forced to read between the lines to see why he proposes entrenchment.

Let me elucidate that point. As I noted at the beginning of this section, it is easy to regard

entrenchment with tradition of use *simpliciter*. This supports my inclination not to take the restriction to projective uses too seriously (see fn. 13).

Goodman's arguments against positionality, for instance, as framed in a justificatory mode: grue-speakers would be able to give the same justification for the projectibility of *their* predicates that we are attempting to give for ours, thus undermining any justificatory force we were hoping to derive from predicate positionality. However, Goodman's equation of demarcation and justification undermines this. The mere fact that 'green' would be positional if we started from the 'grue' language is not pertinent to whether *we* can demarcate the projectible predicates by appeal to positionality. We do not speak the 'grue' language, after all. (Indeed, if positionality is language-relative, why not projectibility as well? So that if we *were* 'grue' speakers, we would regard 'grue' as projectible, and our judgments of positionality and projectibility would still coincide.) So Goodman's arguments in *Fact, Fiction, and Forecast* do not explicitly reveal what is wrong with positionality *qua* demarcation criterion. To be sure, the ensuing appeal to entrenchment does make sense when we have the goal of demarcation in mind. But if we have that goal in mind, it is unclear why positionality itself fails. Why, then, does he appeal to entrenchment instead?

There are two ways I can think of to make sense of this appeal. One is to say that Goodman vacillates on what counts as 'justification' depending on which proposal he is considering. When he considers his own proposal, he treats mere demarcation as adequate, whereas when he considers alternative proposals like positionality, he requires a more robust notion. While this is certainly possible, I am extremely reluctant to accuse a philosopher of Goodman's caliber of such duplicity. Rather, I am inclined to suspect that his argument against alternative proposals has simply not been overtly stated in *Fact, Fiction, and Forecast*. Only in his reply to Barker and Achinstein do we get the explicit statement of the motivation against alternative proposals. That solution is of course hinted at in *Fact, Fiction, and Forecast*: showing

that positionality is language-relative is *close* to showing that the criterion reduces to tradition of use. But the two are not equivalent, and Goodman leaves us to fill in the gap for ourselves. (This, by the way, is why I have referred to Goodman's motivations as "veiled" and "easily misunderstood.")

On this interpretation of Goodman, the main problem with positionality (for instance) as a solution to his riddle is not merely that the grue-speaker could use the same criterion to claim that 'grue' is projectible. Nor is the problem that the grue-speaker's ability to say this undermines our justification for thinking that 'green' and not 'grue' is projectible. Nor, indeed, is the problem that appealing to positionality is circular since it presupposes that 'green' is privileged over 'grue,' which is what we were aiming to underwrite. Rather, the main problem with an appeal to positionality (according to Goodman) is that it ultimately just reduces to an appeal to entrenchment, and therefore such an appeal does not get to the heart of the matter.

This Goodmanian response is generalizable to proposals other than positionality, or 'representable with a single picture.' (Just *how* generalizable it is I will consider shortly.) Suppose we have some particular criterion, φ , that is supposed to demarcate the projectible predicates, or justify their projectibility. Goodman's plan of attack will be to attempt to show that there are nonstandard interpretations of φ according to which it fits the nonprojectible predicates just as well. Only if we interpret φ in the traditional way do we get an acceptable demarcation criterion. And now, since we are appealing to our traditional use of φ to demarcate the projectible predicates, Goodman will suggest that what is really doing the work here is the tradition, not φ itself.

Let us consider some examples of applying this strategy to see its limitations (or lack thereof). Suppose we follow Jackson (1975) and propose that the projectible hypotheses are

exactly the ones that satisfy a counterfactual condition. Hypotheses of the form “All F s are G s” are nonprojectible when all of the F s in the evidence class are also H s (for some H) and “it is known that the F s in the evidence class would not have been G if they had not been H ” (ibid, p. 123). Here H picks out some common predicate of the examined class of F s, such as the fact that they are all examined, that they are all examined *by now*, etc. We can abbreviate this condition by saying that hypotheses of the form “All F s are G s” are projectible when, for any H applying to all of the members of the evidence class, G -ness is *counterfactually robust* with respect to (variations in) H -ness.

How is this intended to rule out the grue hypothesis but not the green one? Well, it will be helpful if we momentarily drop the modified definition of ‘grue’ that I used above, and recall Goodman’s original definition: something is grue iff it is green and first observed before 2020, or blue otherwise. Call this ‘grue_{obs},’ since the time of observation figures in its definition. Jackson intends to use his counterfactual robustness criterion to rule out the hypothesis that all emeralds are grue_{obs}. He does this as follows. We know that all of the emeralds in the evidence class are both observed before 2020 and grue_{obs}. Filling in the schema for the counterfactual condition from above, take F = “emerald,” H = “observed before 2020,” and G = “grue_{obs}.” Now we evaluate whether this hypothesis satisfies the counterfactual condition: is it true that the emeralds in our evidence class would still have been grue_{obs} if they had not been observed before 2020? Obviously not, according to Jackson. This is because we know that emerald color typically does not depend on observation or time of observation. So if the emeralds in our evidence class had not been observed before 2020, they would still have been green, and therefore not grue_{obs}. Since the counterfactual condition is violated, the hypothesis is not projectible. (The same considerations show that the corresponding green hypothesis is not ruled out by the

counterfactual condition: green emeralds in our evidence class would still have been green if they had not been observed before 2020, since again, emerald color tends not to depend on observation or time of observation.)¹⁶

Note that Jackson's proposal is importantly different from proposals like Goodman's that aim to pick out the projectible hypotheses via the projectible predicates. Jackson considers the projectibility of a given hypothesis to depend both on which predicates appear in the hypothesis and on the truth values of certain counterfactuals relevant to the context in which the hypothesis is being evaluated. Grue-ness might not be counterfactually robust with respect to emeralds and the time of observation, but if we find some other context where it is more plausible to suppose that color depends upon the time of observation, then 'grue' may very well figure in a projectible hypothesis.

That being said, we can still apply a variant of Goodman's response to Jackson's proposal. Consider the claims that (1) grue_{obs} is not counterfactually robust with respect to emeralds and time of observation, and (2) green is counterfactually robust with respect to emeralds and time of observation. Suppose someone were skeptical about these claims; how would we substantiate them? Presumably we would need some method of evaluating counterfactuals. Currently the best proposals in that regard involve appealing to possible worlds

¹⁶Observe that the counterfactual condition does not work if we use 'grue' instead of ' grue_{obs} .' For we would need to find some property H of the emeralds in our evidence class with respect to which grue-ness is not counterfactually robust. But what such property is there? It may be tempting to say that the property is "observed before 2020," as before. For this to work, it would have to be the case that whether the emeralds in our evidence class are grue at the time of observation depends on whether they were observed before 2020. While it might seem clear that these emeralds would have been green, not grue, if they had not been observed before 2020, this is essentially just to assume that all emeralds are always green rather than grue. So if we use the modified definition of 'grue' from above, Jackson's counterfactual condition is applicable only if we know the grue hypothesis to be false on independent grounds.

(e.g. Lewis (1973)). According to standard possible worlds semantics for counterfactuals, we evaluate a counterfactual $p \square \rightarrow q$ by (i) imposing a similarity metric on the space of possible worlds, (ii) going to the nearest world where the antecedent p of the given counterfactual holds, and (iii) checking whether the consequent q of the counterfactual also holds there. The counterfactual is true iff q holds in the closest p -world.¹⁷

Now suppose we use such a semantics to evaluate the following two counterfactuals:

1. If the emeralds in our evidence class had not been observed before 2020, then they would still have been grue_{obs} .
2. If the emeralds in our evidence class had not been observed before 2020, then they would still have been green.

To evaluate (1), we go to the nearest world in which the emeralds in our evidence class were not observed before 2020, and check to see whether they are grue_{obs} there. To evaluate (2), we do the same and check to see whether the emeralds are green there. We are thus examining the same world in both cases, so only one of the two counterfactuals can possibly be true.

But which one is true? That will evidently depend on the similarity metric we impose on the space of possible worlds. On our traditional understanding of similarity, the closest (most similar) world where the emeralds in our evidence class were not observed before 2020 will be one in which they are green and not grue_{obs} . Surely, however, there are nonstandard interpretations of ‘similarity’ according to which the most similar world in which the emeralds in our evidence class were not observed before 2020 is one in which those emeralds are grue_{obs} and not green. Indeed, a grue_{obs} -speaker would presumably have exactly that conception of

¹⁷Here I have glossed over a number of the details of the various possible worlds semantics for counterfactuals, including, most notably, Lewis’s rejection of the “Limit Assumption.”

‘similarity.’ So in order to use a possible worlds semantics for counterfactuals in order to secure the truth of counterfactual (2) and the falsity of counterfactual (1), we have to understand ‘similarity’ according to our traditional use of the word.

Here we are in a familiar position. Goodman will now suggest that we appeal to tradition of use *simpliciter* rather than opt for Jackson’s far more circuitous route that involves appealing to the traditional use of ‘similarity’ between possible worlds that figures in the evaluation of counterfactuals. In this way, even something like Jackson’s counterfactual condition can be seen to reduce to tradition of use, i.e. entrenchment.¹⁸

Goodman’s response to Barker and Achinstein therefore looks to be quite generalizable indeed. Goodman is concerned, however, to argue that entrenchment is the *only* demarcation criterion. If that is correct, then it would appear that he thinks his response can show that *any* proposed criterion actually just reduces to entrenchment. In that vein, suppose we tried a drastically more flat-footed demarcation criterion. Suppose I write down a list *L* of exactly the projectible predicates. (Presumably just using a standard dictionary would give me a very good start in this regard.) Now I propose the following demarcation criterion: the projectible predicates are exactly the ones that appear in *L*. Have I not thereby solved the demarcation problem without appealing to entrenchment?

I am not entirely sure how Goodman would respond to such a proposal, but his reply to Barker and Achinstein suggests that he might say the following. Sure, if I understand, by the words listed in *L*, the predicates (or classes) that we traditionally associate with those words, then my list *L* serves as a perfectly acceptable demarcation criterion. The problem is that we have to

¹⁸This reply shares certain similarities with Roskies’s (2008) reply to Jackson. Roskies effectively questions our counterfactual knowledge, suggesting that in the context of the New Riddle, appealing to such knowledge is illegitimate.

assume that the words I wrote down in my list have their traditional extensions. For example, ‘green’ will appear in L . If that word denotes what ‘green’ is traditionally taken to denote, i.e. the class of green things, then we have no problem. But if I defy tradition and understand ‘green,’ in L , to denote the class of grue things, then ‘is a member of L ’ does not serve as a demarcation criterion of exactly the projectible predicates. So we have to assume that the members of L have their traditional extensions. Again, tradition has been smuggled in.

Once again, I stress that I am not entirely sure Goodman would respond this way. All I can say is that his reply to Barker and Achinstein indicates that he is willing to go at least part way toward a response of this kind. Furthermore, since he seems to be committed to entrenchment being the sole demarcation criterion, it would appear that he has to think something like this is correct—otherwise entrenchment would be a nonunique solution given the presence of L .¹⁹

¹⁹Of course, this assumes ‘is a member of L ’ is not ruled out as a demarcation criterion for some other reason, e.g. because it is not general enough to pick out the projectible predicates of descendant languages. That being said, I think Goodman, of all people, should be wary of objecting to demarcation criteria on the grounds that they will fail in the future. Consider this (slightly simplified) construal of the Goodmanian criterion: the hypotheses that are supported, unviolated, unexhausted, and whose predicates are significantly better entrenched than those of any conflicting hypothesis that is also supported, unviolated, and unexhausted, are projectible. Call the hypotheses described in the antecedent the “entrenched” hypotheses; Goodman’s criterion is therefore:

1. The entrenched hypotheses are projectible.

Now consider the following hypothesis, a good companion of Goodman’s ‘gruesome’ generalizations:

2. The entrenched hypotheses are *protoprojectible*,

where a hypothesis is *protoprojectible* at a time t if it is projectible and t is before 2020 or it is unprojectible and t is after 2020. Hypotheses (1) and (2) are conflicting meta-hypotheses about the projectibility of hypotheses; the first of them is Goodman’s demarcation criterion for the projectible hypotheses. Both are supported, unviolated, and unexhausted if either is. But how are we to decide which one is projectible? That is, how do we figure out whether to project Goodman’s demarcation criterion into the future? To follow Goodman and say that (1) is projectible and (2) is not because ‘projectible’ is better entrenched than ‘protoprojectible’ is to completely beg the question here, for in doing so we would simply be appealing to hypothesis

(If one is worried that this treatment of L —reducing it to entrenchment—is somehow too extreme, I am sympathetic, and I will briefly explore that avenue in the conclusion. Given the increasingly bizarre nature of the dialectic, Goodman is of course welcome to jump ship at this point. But he does so at the expense of trivializing a question that he has spent a good deal of time discussing; if L adequately demarcates the projectible predicates, then Goodman’s construal of the New Riddle is disappointingly insignificant.)

There are two points to note about this discussion. First, similarities with the Kripke-Wittgenstein paradox should be increasingly apparent. In both cases we articulate some criterion—either a criterion of sameness of meaning or a criterion that demarcates the projectible predicates—and then go on to problematize that criterion by interpreting it in nonstandard ways. This speaks to the generality of Goodman’s response: so long as the demarcation criterion is given in words (or even demonstratively), it looks like we will be able to interpret the criterion nonstandardly, thus revealing it to be a roundabout appeal to entrenchment. Second, it is curious that there is one criterion which Goodman’s argument *cannot* undermine in this manner, namely entrenchment itself. Suppose I say, “What about nonstandard interpretations of ‘entrenched’? Does the status of entrenchment as a demarcation criterion not depend on our traditional use of the term?” Well if it does, then our criterion is still entrenchment. That does not make the appeal to entrenchment extraneous, as it does the appeal to positionality or counterfactual robustness,

(1): according to (1), the entrenchment of ‘projectible’ counts in favor of the projectibility of (1), but according to (2), the entrenchment of ‘projectible’ counts in favor of the *protoprojectibility* of (1) and against its projectibility. (Note that (2) says nothing about itself because it does not qualify as an entrenched hypothesis.)

So Goodman does not appear to have the resources to establish the projectibility of his own demarcation criterion. This is not a crippling problem for him, but it does preclude his ability to give what is perhaps the natural objection to L as a demarcation criterion, namely that it is not projectible. Rather, it looks like he must rely on something like the argument suggested in the text, i.e. that an appeal to L just reduces to an appeal to entrenchment.

for instance.

Whether or not Goodman's argument works as well as he seems to think it does, it can certainly be seen to provide motivation for his entrenchment solution. I expect that it is with this argument in mind that he says the following in *Problems and Projects*:

[W]e may by now confidently conclude that no general distinction between projectible and non-projectible predicates can be drawn on syntactic or even on semantic grounds. Attempts to distinguish projectible predicates as purely qualitative, or non-projectible ones as time-dependent, for example, have plainly failed. (1972, p. 357)

Note Goodman's emphasis here on our inability to draw any sort of principled distinction even when we rely on the *semantics* of the predicates. Goodman's problem is typically taken to show that no purely *syntactic* rule can characterize the acceptable inductions. As should be no surprise, Goodman takes his problem to extend to semantic considerations as well. Give him a proposed semantic demarcation criterion, and Goodman will turn it into an appeal to entrenchment using a variant of his reply to Barker and Achinstein.

Goodman goes on to draw some interesting conclusions from the syntactic and semantic indistinguishability of the projectible and unprojectible predicates:

Many recent papers look for psychological grounds for our decisions as to what predicates or hypotheses to project. Were such an explanation found, it would not conflict with my treatment of projectibility in terms of entrenchment of predicates but would merely make the initial choices psychologically determinate rather than matters of chance. The difficulties of providing such a psychological explanation, though, are enormous; for any argument that the initial choices of projectible predicates are determined by some non-random operation of the mind requires showing that these predicates are distinguished by some common and independent characteristic that can be correlated with such an operation of the mind. The unavailability of any such characteristic (which would of itself provide a definition of projectibility) is just what gives rise to the riddle and turns our attention to chance and habit. (ibid, pp. 357-8)

This passage contains some critical insights into Goodman's thought. It is tempting to think that even if there is no intrinsic characteristic of the projectible predicates or classes themselves, they at least all have the extrinsic characteristic of being "programmed," as it were, into our cognitive

architecture. Quine suggests something to this extent by the name of an “innate spacing of qualities” in his (1994). The thought is roughly that we must have some sort of innate cognitive dispositions that allow us to pick up on the predicates or classes that we do. More specifically, Quine suggests that we pick out the same classes as each other because we all share roughly the same cognitive architecture, and he speculates further that we pick out the same classes as nature—the ones that are useful in induction—because natural selection has shaped us to do so. (It would, of course, be immensely more advantageous to make predictions based on the naturally relevant classes than on naturally irrelevant ones.)

Goodman thinks this is nonsense. For one thing (as he goes on to note on p. 358), ‘green’ and ‘grue’ would have been equally advantageous until now—the only selective advantages or disadvantages will occur in the future—so natural selection would have no basis on which to discriminate between them. But even that response ignores the crucial point, which is this. To substantiate the claim that we share innate dispositions that determine which classes we go on to use, those classes themselves would have to have some demarcating feature(s) that would allow our dispositions to “latch onto” them. Goodman himself has argued extensively that there are no such demarcating features. Indeed, we have just been rehearsing those arguments above. Without such a feature, all of the dispositions in the world will not give us the distinction we are looking for. “Dispositions, *to do what?* To latch onto *which* classes?”, Goodman will ask.²⁰ Furthermore, if there *were* such a demarcating feature, that itself would serve as the criterion of projectibility rather than any innate dispositions that we possess—remember that Goodman is only interested in codification, after all.

²⁰[W]hat we can’t say, we can’t say, and we can’t whistle it either” (Ramsey, 1931, p. 238). Nor can we encode it in our genes or dispositions. Note that this feature of the classes that we “latch onto” would presumably have to be a natural one, as opposed to the artificial feature of being denoted by a word in *L*. Surely our minds are not innately programmed to latch onto *that* feature.

So Goodman is strongly opposed to the idea that our use of some predicates (classes) rather than others is somehow innately determined. It is determined, rather, only by “chance choices sanctified by habit” (1972, p. 357). In the end this should not be overly surprising, given the immense influence that C. I. Lewis had on Goodman’s thinking.²¹ Lewis’s conceptual pragmatism was noteworthy for a number of reasons, but at least partly so for its distinctively anti-Kantian emphasis on the impermanence of our concepts:

Our concepts are devised with purpose to catch the significant, the subject of meaningful generalization, at whatever level and in whatever way we may. When particular concepts fail, we merely abandon them—through analysis or organization or abstraction, and so on—in favor of corrected ones, which take cognizance of, and include the ground of, our previous failure. (1956, p. 385)

If we replace the flexibility of our concepts here with that of our predicates, then we essentially have Goodman’s own view: a suitably nominalized version of Lewis’s position, free from what Goodman might regard as dubious mental vocabulary like “meaning” and “concepts,” yet maintaining the transient, accidental nature of our categories.²²

I hope to have shown in this section that Goodman views the New Riddle, and his solution to it, somewhat differently than the typical philosopher does. Indeed, careful examination of Goodman’s arguments against alternative proposals and for entrenchment reveals a way of thinking according to which entrenchment is the only plausible candidate for a demarcation criterion. Admittedly, there is a serious question as to whether Goodman’s arguments actually do force us into this way of thinking. I will consider this question briefly later on. For now, I am going to grant Goodman’s view of the problem: the conception of justification

²¹Goodman notes this influence in a variety of places, e.g. in the forewords to *Fact, Fiction, and Forecast* and *Ways of Worldmaking*, and in his memorial piece “Snowflakes and Wastebaskets” (1972, pp. 416-9).

²²Goodman also critiques the notion of innate concepts in his (1967) and (1969), where the Lewisian influence is again quite apparent.

as demarcation, the accidental nature of our categories, the argument pattern that takes him from alternative proposals to entrenchment, and the overall indistinguishability of classes. Granting him all this, I am going to argue that his entrenchment solution still fails.

5 Problematizing Entrenchment

Let us carefully consider what, exactly, entrenchment is supposed to be. It is first and foremost a property of classes, the extensions of predicates. While we can speak of the entrenchment of predicates, to do so is really to speak elliptically of the entrenchment of their extensions. The entrenchment of a class is its history of (projective) uses²³ in our linguistic community; a class is better entrenched the more it has been used.

One way to bring out the problem with entrenchment, *qua* solution to the New Riddle, is to ask how we come to *learn* facts about entrenchment. How do we learn which classes have the longest histories of use in our linguistic community? To learn this, we would have to figure out which class(es) people are denoting when they use a given predicate. So figuring out facts about entrenchment is something that we would have to do in the process of learning a language. Thus it might be thought that while there are undoubtedly complexities associated with this learning process, there could not be anything that is ultimately problematic. We *do* learn languages, after all, even when our cognitive faculties are still inchoate.

The relevant question, however, is whether Goodman's view can account for this. Suppose I ask how I know what class you denote, or have been denoting, with your use of 'green.' In other words, how do I figure out the extension of 'green,' as you use it? Particularly, how do I know you denote the class of green things rather than the class of grue things? My first

²³I will drop the restriction to only projective uses for the remainder of the paper (see fns. 13 and 15). For the purposes of my critique of Goodman's theory, nothing important will come of this.

reaction when I consider this question is to appeal to ostension. I learned how to apply ‘green’ when my parents pointed to a number of green things: a Christmas tree, a patch of well-watered grass, maybe an emerald—eventually I caught on. Now, I certainly “caught on” enough to know that ‘green’ is not applied to blue things. No Christmas trees, grass-patches, or emeralds are blue, after all. But saying that I “caught on” and leaving it at that is insufficient here, for those green things were all grue as well. Any ostensive definition of the class of green things is also an ostensive definition of the class of grue things (at least before 2020), and of innumerable other classes as well. Why should I learn to apply ‘green’ to the class of green things rather than to the class of grue things by your pointing at things that belong to both? Among others, Wittgensteinians should not be surprised by this. It is essentially a restatement of Wittgenstein’s point that “an ostensive definition can be variously interpreted in *any* case” (PI, §28). Recall his reaction to Augustine’s description of a child learning a language ostensively:

Augustine describes the learning of human language as if the child came into a foreign country and did not understand the language of the country; that is, as if he already had a language, only not this one. Or again, as if the child could already *think*, only not yet speak. And “think” would here mean something like “talk to himself.” (ibid, §32)

Ostension is only very minimally helpful in the acquiring of concepts and the learning of language; it *assumes* that we already have the relevant concepts (it gives us a sound to attach to them) rather than teaching us the concepts themselves. So it cannot play the role that we need it to: it cannot mark off the extensions of our terms unambiguously.

Can you do more than ostend in order to communicate to me the class you denote by ‘green’? Well, you could try to *describe* the class you are denoting. But there would be problems there as well, and all too familiar ones at that. Joseph Ullian nicely elucidates the difficulties:

How is it that general terms turn out to be of use for us? Now it is a brute and happy fact that human beings are able to learn. . . how to use such terms. We, and with us other creatures, are able to achieve *recognition* of what the myth makers have reified as

properties...But you can explain what this ability consists in only up to a point. Any explanation of how we learn to apply a specified general term will itself make use of general terms, which—far from being a defect—is reflective of a basic feature of language and a *sine qua non* of it (the need for general terms or their equivalents in saying *anything*). You can point out that we see a few examples of blue things and we catch on, if judiciously taught, how and when to apply ‘blue’, but what more can you say? You certainly explain nothing if you say that we learn to apply ‘blue’ to blue things. And you only push the problem further back if you say that we learn to apply ‘blue’ to things sufficiently similar in color to the initial specimens. For in what does similarity in color consist? Somewhere you must terminate the attempt at explanation. . . I think there is a natural temptation to give too much credit to nature for all this, as if objects and simple properties appeared for us already labeled, and our only job was to read the labels (or as if they carried membership cards for exclusive clubs known as *natural kinds*). Concepts and the employment of general terms depend upon classifications, and, as I like to put it, one class is as good as another as far as logic is concerned. (1961, p. 736, italics in original).

So if you try to describe the extension of your predicates, I can always worry that I am using a different interpretation of some of the words of your description. I can worry, for example, that we might be using different notions of ‘similarity in color,’ in the case of color concepts, just like I can worry that we might be using different notions of ‘counting’ in the case of addition. You can tell me that green things will still look green, and not blue, after 2020, but how do I know you do not understand ‘green’ and ‘blue’ here as ‘grue’ and ‘bleen’? For then you would say the same thing. And you can tell me that the classes you denote are representable with a single picture, but how do I know what notion of ‘representation’ you have in mind? Remember, according to Goodman there is ultimately *no semantic distinction* here. Any description of the extension of a given predicate is going to involve other predicates, and so will only push the problem back onto them. So it looks like both ostension and description are of no help.

Neither pointing nor speaking will fix, for me, the class that you denote by your use of ‘green.’ Still, it might seem that I can get a grip on the extensions of your predicates by appealing to our shared cognitive architecture. Roughly: your ‘green’ must pick out the same class as my ‘green’ because we share the same conceptualization abilities, and we are naturally

inclined to go on in the same way once we have seen ‘green’ applied to a few samples. This is perhaps a reasonable response, but can Goodman make it? Evidently not, given the passage of his that I quoted above (1972, pp. 357-8). To appeal to shared innate dispositions to pick out the same classes would be to presuppose that the classes themselves have some demarcating feature that our dispositions could latch onto, which is exactly what Goodman has argued against.

It now begins to look as though Goodman has precluded any way of securing knowledge of facts about entrenchment. Indeed, suppose there *were* some method by which you could fix the extensions of your predicates such that I could learn them unambiguously, whether by pointing or speaking or appealing to our shared cognitive capabilities. Then there would be some demarcation criterion *distinct from entrenchment* that distinguishes ‘green’ from ‘grue.’ If the extensions were fixed by description or ostension, then that description, or a description of the relevant act of ostension, would serve as the demarcation criterion. If the extensions were fixed by appeal to our shared dispositions, then by Goodman’s own argument there would have to be some distinguishing feature of the classes themselves that those dispositions could latch onto.

Here is another way to bring out the problem with entrenchment. Recall the list *L* of projectible predicates that I proposed as a demarcation criterion above. There I said that if Goodman wants to be consistent with his response to Barker and Achinstein, he should say that whether *L* serves as a demarcation criterion depends on whether we interpret the predicates appearing in the list in the standard way. Now suppose I just speak the list, rather than writing it down. Would my spoken version of *L* (call it ‘*L_S*’) serve as an acceptable demarcation criterion? If Goodman responds consistently, he will say that it depends on whether we understand the words I have spoken to have their traditional extensions. But now we are in dangerous territory, for ‘traditional extension’ here just means ‘what I, and the members of my linguistic community,

traditionally use the words to denote.’ And if Goodman can question what I am denoting in the case of L_S , then what is to stop him from questioning what I, and others, have been denoting on other occasions by the use of those words? That is, how do we tell what the traditional extensions are?

Of course, as soon as we start wondering how (or whether) we know what the traditional extensions are, we likewise have to start wondering whether there are such things. Take the predicate ‘green’, for example. If we can use Goodman’s arguments to work ourselves into a state where we are unsure about what the traditional extension of ‘green’ is, then we will also have to wonder whether everyone, or the majority of people, have been denoting the *same* class by the use of the term. Maybe John has been denoting the class of grue things, Matt the class of gred things, and Marc the class of grelow things. As of yet, nothing has ruled out such a state of affairs. It is tempting here to say that there is a remarkable amount of evidence suggesting that we *do* all use the term with the same denotation. For what else could explain our massive agreement in its application thus far? But this would be just like saying that we have a remarkable amount of evidence that emeralds are all green, because they have all been green thus far. Such an argument might have traction in a context where we are not concerned with the validity of induction, but here it is merely question-begging.

It is important to notice that nothing in my argument against entrenchment relies on there actually being no uniquely well-entrenched classes. Perhaps, by the purest of chances, when people have used ‘green,’ the overwhelming majority have used it to denote the class of things that I think of as green. If Goodman’s arguments are correct, such a state of affairs is extremely unlikely (think of all the theoretically available alternatives!), but conceivable. So maybe there are uniquely well-entrenched classes for some, or even all, of our predicates. Still, the problem is

that this would be of no use to us. Remember, Goodman is looking for an *informative* demarcation criterion. This was the lesson of the above discussion of ‘projectible,’ which seemed to serve as an adequate demarcation criterion on a naïve reading. What I have been arguing is that his own arguments show that entrenchment cannot be informative.

Let me briefly allay a concern one might have—that I have been appealing to a more acute kind of skepticism than Goodman is concerned to dispel. My argument against entrenchment is not meant to rely on something like the inverted spectrum argument. I do not mean to be merely pointing out that your green could be my blue, for example. That is, I do not mean to be appealing to the independent (and far more antiquated) skeptical worry that we might mean different things by the use of our words, so that appealing to entrenchment is problematic because it runs up against this older worry. What I am suggesting is that the entrenchment solution encounters the *same* problems that it was designed to solve. In order for an appeal to entrenchment to work as a solution to the New Riddle, I do not have to know that we *mean* the same thing by ‘green,’ but I have to have good reason to think that we apply it in the same cases. Otherwise I do not have a grip on the extensions of your predicates, and thus I do not have a grip on facts about entrenchment. So the concern is this: how do I know you apply ‘green’ in all the same situations in which I apply it? Granted, I know that you have not, in any past moment, applied ‘blue’ to something that was green at that moment. But I do not know that you apply ‘green’ to green things as opposed to grue things, or gred things, or grellow things, etc. And the latter cases are exactly the problematic cases that we were trying to decide among. One might put the problem this way: Goodman assumes that we can have a better grip on the classes used by other humans than on the classes used by nature, but his arguments undermine knowledge of both equally. (Indeed, in both cases we are projecting regularities in our experience. Whence the

relevant difference?) Outside of a Goodmanian context, one could confirm that you and I apply ‘green’ in the same cases by pointing to past instances where we have agreed in application conditions, or to some shared innate concepts that we possess as humans. Both of those are ruled out if we have Goodmanian concerns in the picture. Past coextension in our applications of ‘green’ does nothing to guarantee future coextension unless certain classes are privileged.

Having argued that the entrenchment solution falls victim to the same considerations that give rise to the New Riddle, I want to consider some possible responses on behalf of Goodman. The most obvious response is that something must have gone wrong if what we have shown is essentially that it is impossible to learn a language, or (to put the same issue differently) to know that you and I are speaking the same language. Perhaps the correct reaction to the above argument against entrenchment is to reject some of the assumptions that led to it rather than to accept such an extreme conclusion. Aside from noting that this reaction would likewise undermine the motivation for the New Riddle itself (and thus would not really be “on behalf of Goodman”), I want to hold off on addressing this response until later. This is because I am going to push on the conclusions I have drawn so far, so I would prefer to wait to consider the “Modus Ponens or Modus Tollens?” question until the full conclusion is in view.

So what could Goodman say in response to my argument? Perhaps it could be claimed that he need not say anything; he has already said it at the beginning of *Fact, Fiction, and Forecast*:

A philosophic problem is a call to provide an adequate explanation in terms of an acceptable basis. If we are ready to tolerate everything as understood, there is nothing left to explain; while if we sourly refuse to take anything, even tentatively, as clear, no explanation can be given. What intrigues us as a problem, and what will satisfy us as a solution, will depend upon the line we draw between what is already clear and what needs to be clarified. Yet I am afraid that we are nowhere near having any sound general principle for drawing this line. (1983, p. 31)

Goodman may indeed respond this way. Surely we can take it as an item of knowledge that we all speak the same language, i.e. that we all denote the same classes with our predicates. Surely, that is, facts about entrenchment are on the “clear” side of the line he mentions dividing the clear and the unclear. So for me to argue that facts about which classes other humans use are no “clearer” than facts about which classes nature uses is simply for me to draw that line far too narrowly.

I would parry as follows. Firstly, if we have no sound general principle for deciding where to draw the line, then at best Goodman’s rebuttal boils down to the claim that he wants to draw the line differently than I do. Fair enough, I suppose, but now we are no longer *debating* as much as we are stating our preferences regarding the clarity of certain notions. But secondly, and more importantly, even if we do not have a sound and general principle for deciding where to draw the line between clarity and unclarity, surely we have some principle that can tell us whether or not that line itself has been drawn in a disciplined manner. And surely that principle would tell us that if Goodman responds in this way, he has drawn a very tortuous line. If the same arguments that motivate the worry about projection with respect to nature also motivate worries about projection with respect to which classes our fellow humans denote, then how can we justifiably appeal to knowledge of the latter to secure knowledge of the former? Admittedly, such arguments may not tell us conclusively whether to regard both items as clear or both as unclear, but I submit that they do tell us that we should regard them both in the same manner.

Another way for Goodman to respond is the following. What he claimed to be denying was that there is no *general* distinction between projectible and unprojectible predicates. That is, there is no single criterion by which we can separate all of the projectible predicates from all of the unprojectible ones. This is different from saying that there is no distinction between two

particular predicates like ‘green’ and ‘grue.’ That there is no general distinction between the projectible and the unprojectible predicates does not imply that there is no particular distinction between a given pair of predicates. As long as we can draw these more particular distinctions, then we can get a grip on facts about entrenchment. The presence of such a distinction, for example, would allow me to figure out that you are denoting the class of green things rather than the class of grue things with your use of ‘green.’

The problem with this response is that it does not square with his other comments. Granted, Goodman does sometimes state his concern as one of finding a general demarcation criterion, but he argues for the lack of such a criterion by arguing that given proposals will not work to distinguish *particular* pairs of predicates. As an example, consider his treatment of the positionality criterion. I noted above (fn. 11) that he does *not* object to the positionality criterion, either in *Fact, Fiction, and Forecast* or in his reply to Barker and Achinstein, by pointing to problems of generalizing it. That is, he does not point out that the the class of nonpositional predicates likely cross-cuts the class of projectible predicates, so that if we use it as a demarcation criterion, it will sometimes get the wrong results—it will deem unprojectible predicates projectible, or vice versa. Rather, his focus in both cases is on the fact that positionality does not even establish a distinction between ‘green’ and ‘grue’ (unless we read it as a covert appeal to entrenchment).

Ullian gives a similar reply to Goodman in *his* (1961b) response to Barker and Achinstein, which he describes as being “encouraged” by Goodman. Ullian argues that our judgments of whether a given color “matches” the pigment we use to represent it will differ from that of the grue-speaker’s. We could emphasize this, he suggests, by saying that the grue-speaker has a notion of ‘sameness of schmolor’ rather than ‘sameness of color.’ Thus the grue-speaker will use

grue (i.e. green) pigment to represent grue things, even post-2020 grue things. Both Goodman and Ullian, therefore, seem determined to undermine any way of distinguishing between even this *particular* pair of predicates. Indeed, Goodman's extensive reliance on this one particular example ('green' vs. 'grue') indicates that he thinks the primary problem can be posed as one of distinguishing particular pairs of predicates rather than distinguishing classes of them.

It is also worth noting that if the problem were primarily one of the *generality* of the demarcation criterion, then the list *L* would appear to be perfectly adequate. Indeed, that criterion is *designed* to be thoroughly general, at least insofar as the generality of a demarcation criterion refers to its ability to uniformly separate all of the projectible predicates from the unprojectible ones. So if Goodman's problem is not trivially solved by the posit of *L*, the problem cannot be primarily one of the generality of the demarcation criterion.

A final response on behalf of Goodman would be to appeal to *personal* entrenchment, by which I mean *my* history of use of a given predicate, rather than the history of use in my linguistic community as a whole. There are two obvious ways Goodman could deploy personal entrenchment: either (1) use it to bootstrap our way up to communal entrenchment or (2) directly treat personal entrenchment as our demarcation criterion. I will consider these proposals in turn.

The first proposal is easy to dispatch. To use personal entrenchment as a guide to communal entrenchment is just to presuppose that I have something in common with others in my community that leads us to pick out the same classes. We have already seen the problem here: Goodman himself rules out this possibility in the passage from *Problems and Projects* above. As I emphasized there, both that passage and his sympathies with Lewisian conceptual pragmatism undermine the first proposal.

Evaluating the second proposal leads to more interesting conclusions. Here we are

attempting to use my past history of use to demarcate the projectible predicates; the projectible ones will be those that I have used the most, or those that are coextensive with the ones that I have used the most. This would seem to be the natural move for Goodman to make at this point. Sure, maybe I do not ultimately have a grip on which classes other people have been denoting when they use predicates like ‘green.’ Nevertheless, I undoubtedly have access to the classes that *I* have been denoting by my use of a given predicate, and that can serve just as well to demarcate the projectible predicates. This response preserves the spirit of Goodman’s solution, even maintaining his Humean emphasis on habit as the source of our judgments about projectibility.²⁴

If one does not look too carefully, this fallback can appear perfectly acceptable. However, similar considerations to those that undermined communal entrenchment can also compromise this response. Again we will appeal to the indistinguishability of ‘green’ and ‘grue’ to problematize the proposal, but this time the appeal will rely on metaphysical concerns, rather than epistemological ones. Consider: if Goodman is right that there are no features of the predicates themselves by which they can be distinguished, then what could make it the case that I was using one and not the other in the past? Goodman essentially takes it as given that the two predicates are different, i.e. have different extensions, but then goes on to argue against every way in which they might differ, aside, of course, from entrenchment. But entrenchment alone cannot make a difference where there previously was none. For two predicates to differ in their histories of use, they must first *differ*. (History of use *of what?*) The same considerations clearly apply to personal entrenchment. If there are no features that distinguish them, then what makes

²⁴“I submit that the judgment of projectibility has derived from the habitual projection, rather than the habitual projection from the judgment of projectibility. The reason why only the right predicates happen so luckily to have become well entrenched is just that the well entrenched predicates have thereby become the right ones” (1983, p. 98).

them different predicates?

Something like Leibniz's Law of the identity of indiscernibles is operating in the background here: if two things share all the same properties, then they are identical. I do not mean to commit myself to the metaphysical necessity of such a principle in general, but in the case of my use of a given predicate or concept, the principle seems perfectly apt. To deny it would be to maintain that I can use one predicate (or concept) rather than another even if there is no difference between the two. Such a proposal leads us into treacherous territory, both metaphysically and epistemologically. If it were correct, what would determine which predicate I am using? My intention, maybe? My intention to do *what*, exactly? And if it were correct, how would I know that all of my past uses of 'green' really were uses of the same predicate? To deny Leibniz's Law with respect to concepts or predicates is metaphysically incoherent and epistemologically hazardous. If there is such a thing as using a given predicate or concept, it must be subject to this conceptual version of Leibniz's Law.

So if Goodman appeals to personal entrenchment to distinguish 'green' from 'grue,' the response is that his arguments undermine even the possibility that I can use one and not the other. With due respect to Zabludowski (1977), *quod perit tandem perit*. To put it figuratively, if there is no feature by which I can latch onto one class rather than another, then there is likewise no way to hold onto the same class over time, or even to pick out one class rather than another at a single time. This leaves us in a very familiar position. It is essentially the same problem that Kripke-Wittgenstein's skeptic has posed with respect to 'plus' and 'quus.' Of course, that problem was posed in terms of the question of how I could *mean* one and not the other, whereas this problem is now posed in terms of the question of how I can *use* one and not the other. I am inclined to attribute this difference merely to Goodman's reticence to make extensive use of the

word ‘meaning’ and its cognates. That is, it is difficult to see how this would be an important difference.

It is worth noting that this problem with personal entrenchment undermines another response that Goodman could try to make to my argument. Ostensibly, Goodman can just claim that, in posing the New Riddle, he is presupposing that something like the Kripke-Wittgenstein paradox can be solved, just as he is (justifiably) presupposing that external world skepticism can be dealt with. If he can presuppose the latter, why not the former as well? Answer: because Goodman’s arguments generate worries not just about demarcating the projectible predicates from the unprojectible ones, but about demarcating my use of ‘green’ from my use of ‘grue,’ ‘plus’ from ‘quus,’ etc. Goodman’s demarcation problem is a special case of Kripke-Wittgenstein’s. Thus for Goodman to pose the question of how to demarcate the projectible from the unprojectible predicates while presupposing that the Kripke-Wittgenstein paradox can be solved would be analogous to presupposing that external world skepticism can be solved, and then asking what solves external world skepticism on the weekdays. The appropriate answer would just be: whatever solved the general problem in the first place.

If what I have said is right, not only does the appeal to personal entrenchment fail, but also there is something misleading about the way the New Riddle has been posed. Granted, following Goodman’s arguments to their full conclusions *does* lead us to a puzzle about how to demarcate the projectible predicates (hypotheses) from the unprojectible ones. However, this puzzle is really just a special case of the predicament in which his arguments leave us. Not only are we unable to demarcate the projectible predicates from their gruesome cousins, but we are also unable to distinguish between our meaning (using) green as opposed to grue, or plus as opposed to quus. We are accustomed to viewing the New Riddle as one in which the evidence

statements are taken both as given and as perfectly understood, and our job is to figure out how to form predictions on the basis of those statements. Now I am suggesting that Goodman's arguments undermine even what we thought was given; the statements themselves have been problematized.

In *Wittgenstein on Rules and Private Language*, Kripke anticipates some of the similarities between Goodman and Wittgenstein that I have been stressing here. For example, in footnote 44 on page 58, he says that "Goodman's discussion of [the New Riddle] seems to presuppose that the extension of each predicate ('green', 'grue'), etc., is known, and that this question does not itself get entangled in the 'new riddle of induction'." I have been arguing that the notion of the extension of a predicate does indeed get entangled in the New Riddle; extensions themselves, just like meanings, are problematized by Goodmanian considerations. Kripke goes on to suggest that "serious consideration of Goodman's problem...may prove impossible without consideration of Wittgenstein's" (p. 59).²⁵ It should be abundantly clear from the above considerations that I agree with him.

6 Conclusion

By way of concluding, I want to offer both a clarification and some speculation.

Regarding the clarification: nothing I have said precludes one's ability to view the New Riddle differently than Goodman. That is, we can assume that the Kripke-Wittgenstein paradox has been solved, or at least justifiably set-aside, and we can take the New Riddle to be asking a more substantively justificatory question than Goodman suggests. If the Kripke-Wittgenstein paradox *is* solved, however, then presumably this will involve pointing to some fact that makes it the case that I can mean plus rather than quus, and green rather than grue. So if it is solved, it

²⁵See also his footnote 46 on page 59.

looks like some way of distinguishing between predicates (or their associated classes) will have to be found. Thus the demarcation problem would be solved, and there would still be a “riddle” here only if we employed a more robust notion of justification than Goodman. Much of the literature on the New Riddle can be viewed as adopting this perspective: assume there is a way to solve the demarcation problem, and ask the further question of why we are *justified* in projecting the predicates that are so demarcated. Though I think this is a perfectly fine question to ask, I also think we should be paying more attention to the demarcation problem that it presupposes can be solved. It is not so clear to me that it can be.

As for the speculation, I want to briefly consider whether Goodman’s arguments actually do force us into the Kripke-Wittgenstein paradox. As I argued above, they do so if they establish the points that Goodman takes them to establish. Now, do they actually establish those points?

There are perhaps a few places where it might seem that Goodman has moved too quickly. The most noteworthy is his claim that there are no distinguishing features between the projectible and the unprojectible predicates. I suggested above that he says this with something like his response to Barker and Achinstein in mind. Namely, any alternative demarcation criterion can be turned in a roundabout appeal to entrenchment by considering nonstandard interpretations. Does such an argument really establish that there are no features of ‘green,’ for example, that are not also features of ‘grue’?

It is not obvious to me that it does. Indeed, would Goodman’s argument not apply more generally, allowing us to claim of any two objects whatsoever that they share all the same features? If so, surely something has gone wrong. But let us maintain our focus on the narrower task of distinguishing just between predicates, e.g. ‘green’ and ‘grue.’ What Goodman *can* do here is ask us to give him some feature of one that the other lacks. When we do this, he will twist

it into a feature of the other by some nonstandard interpretation of that feature. Nevertheless, it might be thought that what this maneuver shows is just that there is nothing we can *say* about one that we cannot also say about the other. That does not necessarily imply that the two predicates share all of the same features; we might still be able to maintain a difference between the two if we were content to let it be an ineffable difference. Then the skeptical worries about how I can use (or mean) one and not the other, if there is indeed no difference between them, would fall away. There *would* be a difference between them, albeit an inarticulable one. If so, then that difference itself could serve as traction for some shared dispositions to pick out the same classes.

I doubt that Goodman would be on-board with this. Particularly, he would think that talk of the ineffable is misguided. Consider the following excerpts from the first essay of *Ways of Worldmaking*:

1. "If I ask about the world, you can offer to tell me how it is under one or more frames of reference; but if I insist that you tell me how it is apart from all frames, what can you say? We are confined to ways of describing whatever is described. Our universe, so to speak, consists of these ways rather than of a world or of worlds" (p. 3).
2. "Predicates, pictures, other labels, schemata, survive want of application, but content vanishes without form. We can have words without a world, but no world without words or other symbols" (p. 6).

If I understand these passages correctly, Goodman thinks the notion of the ineffable is confused. "*No world without words or symbols.*" If so, then our inability to articulate a difference between 'green' and 'grue' does indeed imply that they share all the same features, just as Goodman says.

I will not here enter into a discussion of the plausibility of the ineffable, but the

coherence of that notion certainly does seem to be relevant to whether Goodman's arguments force us into the Kripke-Wittgenstein paradox. I should only like to point out, in passing, how odd it is that this issue has shown up here, and to note that it will likely be relevant to deciding whether we ought to follow Goodman's arguments to their conclusions, or take the absurdity of the result to indicate an absurdity in the starting point. In other words, I do not think the *mere* absurdity of the KW paradox should lead us to reject the Goodmanian arguments that lead us there. If something has gone wrong, we should be able to locate it. And that necessitates a search for which I do not here have the space.

REFERENCES

- Barker, S. F. and Peter Achinstein. "On the New Riddle of Induction." *The Philosophical Review*. 69.4 (1960): 511-22.
- Goodman, Nelson. "A Query on Confirmation." *The Journal of Philosophy*. 43.14 (1946): 383-85.
- Goodman, Nelson. *Fact, Fiction, and Forecast*. Fourth Edition. Cambridge: Harvard UP, 1983.
- Goodman, Nelson. "Positionality and Pictures." *The Philosophical Review*. 69.4 (1960): 523-25.
- Goodman, Nelson. *Problems and Projects*. Indianapolis: Bobbs-Merrill, 1972.
- Goodman, Nelson. "The Emperor's New Ideas." In *Language and Philosophy* Ed. Sydney Hook. New York: New York UP, 1969. 138-42.
- Goodman, Nelson. "The Epistemological Argument." *Synthese*. 17.1 (1967): 23-28.
- Goodman, Nelson. *Ways of Worldmaking*. Indianapolis: Hackett, 1978.
- Goodman, Nelson and Catherine Elgin. *Reconceptions in Philosophy and Other Arts and Sciences*. Indianapolis: Hackett, 1988.
- Hacking, Ian. "Entrenchment." In *Grue!: The New Riddle of Induction* Ed. Douglas F. Stalker. Chicago: Open Court, 1994. 193-223.
- Jackson, Frank. "Grue." *The Journal of Philosophy*. 72.5 (1975): 113-31.
- Kripke, Saul. *Wittgenstein on Rules and Private Language*. Cambridge: Harvard UP, 1982.
- Lewis, C. I. *Mind and the World-Order*. New York: Dover, 1956.
- Lewis, David. *Counterfactuals*. Malden: Blackwell, 1973.
- Quine, W. V. "Natural Kinds." In *Grue!: The New Riddle of Induction* Ed. Douglas F. Stalker. Chicago: Open Court, 1994. 41-56.
- Ramsey, F. P. "General Propositions and Causality." In *F. P. Ramsey: The Foundations of Mathematics*. Ed. R. B. Braithwaite. London: Routledge and Kegan Paul, 1931.
- Roskies, Adina. "Robustness and the New Riddle Revived." *Ratio*. 21.2 (2008): 218-30.
- Schwartz, Robert, Israel Scheffler, and Nelson Goodman. "An Improvement in the Theory of Projectibility." *The Journal of Philosophy*. 67.18 (1970): 605-8.

- Stalker, Douglas F., Ed. *Grue!: The New Riddle of Induction*. Chicago: Open Court, 1994.
- Ullian, Joseph. "Luck, License, & Lingo." *The Journal of Philosophy*. 58.23 (1961a): 731-38.
- Ullian, Joseph. "More on 'Grue' and Grue." *The Philosophical Review*. 70.3 (1961b): 386-89.
- Wittgenstein, Ludwig. *Philosophical Investigations*. Trans. G. E. M. Anscombe, P. M. S. Hacker, and Joachim Schulte. Malden, MA: Wiley-Blackwell, 2009.
- Zabludowski, Andrzej. "Concerning a Fiction About How Facts Are Forecast." *The Journal of Philosophy*. 71.4 (1974): 97-112.
- Zabludowski, Andrzej. "Revised but Not Secured: A Reply to Ullian." *Journal of Philosophy*. 79.6 (1982): 329-32.
- Zabludowski, Andrzej. "Quod Periit, Periit." *The Journal of Philosophy*. 74.9 (1977): 541-52.