The Disjunctive Riddle and the Grue-Paradox

Wolfgang FREITAG*

ABSTRACT

The paper explores the *disjunctive riddle* for induction: If we know the sample *K*'s to be *P*, we also know that they are *P* or *F* (for arbitrary *F*). Assuming that we also know that the future *K*'s are non-*P*, we can conclude that they are *F*, if only we can inductively infer the evidentially supported *P*-or-*F* hypothesis. Yet this is absurd. We cannot predict that future *K*'s are *F* based on the knowledge that the samples, and only they, are *P*. The ensuing challenge is to account for the unprojectibility of the disjunctive hypothesis. I provide an explanation in terms of *epistemic dependence*. The *P*-or-*F* hypothesis is unprojectible because the evidence supporting it depends epistemically on the evidence for the defeated *P*-hypothesis. The paper also shows that the disjunctive riddle covers the essence of Goodman's infamous grue-problem, which therefore can be resolved by the same means: In contrast to the green-hypothesis, the grue-hypothesis is unprojectible because the grue-evidence depends on the evidence for a defeated hypothesis.

Suppose we have randomly drawn (without replacement) 99 balls from an urn containing 100 balls and found that they are made of plastic. This gives us excellent evidence for the hypothesis that the 100th, and final, ball is plastic, too. If, on the other hand, we were to receive the additional information that there had been a wooden ball in the urn before the drawing, we would no more be willing to accept the general hypothesis that all balls, the final one included, are plastic. Our knowledge that the final ball is made of wood functions as a defeater for the otherwise perfectly confirmed hypothesis that all balls are plastic.

Now, let us imagine that someone were to exploit this situation in the following way. "Granted", she says, "we cannot predict that the final ball is made of plastic. Nevertheless, we can infer a whole lot of other things. In particular, we may legitimately conclude that it is blue, radioactive, and that it will be kissed by the Queen." Upon being challenged, she justifies her astonishing claims by reference to the following type of reasoning: "We have found the samples to be plastic and, by simple disjunctive weakening, to be plastic or *F*, for arbitrary *F*. Given that all samples are plastic or *F*, there seems to be nothing whatsoever precluding the inductive inference to the hypothesis that all balls, the final one included, are plastic or *F*: there is no background knowledge that conflicts with the disjunctive hypothesis.[1] However, given the perfectly confirmed general hypothesis that all balls are plastic or *F*, we now proceed to deduce, knowing that the final ball is not plastic, that it is *F*! Since we may stipulate that *F* vicariously stands for the properties of, say, being blue, radioactive, and attractive to the Queen, we are able to infer that the final ball has all of these properties!"

Obviously this argument is preposterous. We cannot derive any interesting conclusion from the mere fact that the 99 samples are plastic and that the 100th ball is made of wood. To locate the problem, observe that there is nothing wrong with the

* University of Freiburg, University College Freiburg/Department of Philosophy, Bertoldstr. 17, D-79085 Freiburg; Email: wolfgang.freitag@philosophie.uni-freiburg.de

[1] I here assume that we do *not* know that the final ball is not *F*.

two *deductive* steps involved. Disjunctive weakening is impeccable and so is the employment of a disjunctive syllogism. The fault must reside in the inductive inference alone: the disjunctive hypothesis that all balls are plastic or *F*, although supported by positive instances, is not, to borrow Nelson Goodman's phrase (Goodman 1983, 74), genuinely confirmed. A solution to the *disjunctive riddle*, as I call this problem, explains why the projection of a disjunctive weakening (e.g., "plastic or *F*") is unprojectible, if that of the predicate itself ("plastic") is.

In this paper, I propose a solution in terms of 'derivative defeat'. A hypothesis is *derivatively defeated* if the inductive evidence supporting it *epistemically depends* on the inductive evidence for a (directly) defeated hypothesis. The projection of "plastic or *F*" is derivatively defeated, because the hypothesis of the last ball's being plastic is directly defeated by our knowledge that it is wooden. Like the riddle itself, the solution is shamefully simple, but it suggests that there is something seriously amiss in any theory of confirmation that exclusively focuses on *what* we believe and ignores *why* we believe what we believe. The present paper emphasizes the relevance of epistemic dependence for inductive inferences. So I consider it to be of importance beyond the narrow concern with some petty puzzle of confirmation.

Furthermore, our problem is not quite as insignificant as it may appear at the outset. Firstly, observe that the disjunctive riddle can be constructed on the basis of *any* predicate that applies to all the samples and only to them. It can, for instance, be framed around the predicate "sampled" (instead of "plastic"), i.e., a predicate that is trivially true of the samples and of them exclusively – no matter what the samples are. As a consequence, the disjunctive riddle arises for any induction scenario whatsoever and hence constitutes a pressing problem for confirmation theory in general. Secondly, while the disjunctive riddle has been largely neglected in just the form I have given it, it is very familiar in a slightly different version, widely known as 'Goodman's paradox'. As will be shown below, "grue" is but a logical strengthening of "sampled or blue". The solution to the disjunctive riddle thus functions as a solution also to the famous grue-paradox: the grue-hypothesis turns out to be derivatively defeated. In contrast to many prominent alternative suggestions,[2] I therefore propose that unprojectibility is not an intrinsic feature of ill-conceived predicates, but a consequence of the epistemic conditions created by them.

The paper systematically develops this line of argument and thereby some ideas recently sketched in Freitag 2015. I begin with an analysis of simple cases of defeat (sect. 1). After describing the disjunctive riddle (sect. 2), I develop the solution in terms of epistemic dependence and derivative defeat (sect. 3). Section 4 transfers this type of solution to Goodman's paradox and claims that the grue-hypothesis is derivatively defeated. It also considers some objections.

1. Discriminating predicates and direct defeat

For reasons of simplicity, I confine myself to enumerative induction. Let an induction set *I* be the ordered pair of two disjoint sets, the sample set I_α, consisting of $\alpha_1, \alpha_2, \ldots, \alpha_n$, and the (possibly infinite) target set I_β that comprises the objects onto which to project, i.e., $\beta_1, \beta_2, \ldots$, such that all α's and β's belong to a single kind of objects *K* (e.g., balls from/in the urn, men, emeralds). Inductions of the form here

[2] For a collection of standard approaches to Goodman's paradox, see Stalker 1994.

considered are inferences from the proposition that all α's have property P (or some logically stronger property P^*) to the general hypothesis that all K's, the β's included, are P. In this way we may predict, for example, the color of future balls (the β's) on the basis of the color of past ones (the α's). I will then also say that "P" is *projected* from the α's to the β's. I assume that I_α and I_β are nonempty: there are always sample K's, and K's onto which to project. In this way we avoid that a projection is trivially confirmed by a lack of α's and that it is exhausted because there are no β's.

For the purposes of simple exposition, I will reserve the term "inductive evidence proposition" to statements of the form "$P\alpha_i$", i.e., statements about elements of I_α alone. Inductive P-evidence is the totality of inductive evidence propositions of the form "$P\alpha_i$". That all the randomly drawn balls are made of plastic constitutes the inductive evidence with regard to the material constitution of the balls in the urn case. And that the final ball is *not* made of plastic is part of the total evidence alone.

Call a predicate P (*epistemically*) *discriminating* (with respect to induction set I and subject S[3]) if and only if the epistemic agent S knows (i) that the elements of I_α are P and (ii) that the elements of I_β are not P. By the very definition of the term, the projection of a discriminating predicate P is fully supported by the inductive evidence: all objects in the sample set, all α's, are known by S to be P [by condition (i)].[4] Nevertheless, S must not project P since there is a *defeater*: S knows that the β's are *not* P [by condition (ii)] and hence that the general hypothesis ("All K's are P") is false. Discriminating predicates are such that it is known by S that their extension covers all the samples and only them. They therefore reflect the expectations that are raised by inductive evidence and belied by the defeater:[5] The known fact that all 99 balls randomly drawn from the urn are plastic yields perfect inductive support for the hypothesis that the remaining ball is plastic too, whereas the additional information that the urn contains a wooden ball blocks that prediction. The air of paradox in this case (if there has ever been any) dissolves once we are in command of the notion of a discriminating predicate and hence of a defeater. A theory of confirmation meant to deal with such cases requires no more than a footnote to the effect that positive instances allow for predictions only in the absence of conflicting background knowledge.

Importantly, to be discriminating is not an *intrinsic* characteristic of a given predicate. It is a matter of S's epistemic position in relation to the α's and the β's, respectively. If none of our balls had ever been examined for their material constitution, inductive support for the general statement that all balls are synthetic would be missing (in violation of condition (i)). And if there were no undermining background knowledge, condition (ii) would not be satisfied and the hypothesis in

[3] I will henceforth omit explicit reference to induction sets and subjects if context allows.

[4] While I have been referring to inductive evidence, and the defeater, as kinds of *knowledge*, we must keep in mind that other types of epistemic or doxastic attitudes, e.g., belief or acceptance, would do as well. What is important is not the specific epistemic attitude in operation, but that we deal with an exclusively epistemic or doxastic phenomenon in the first place. The projection of the hypothesis that all K's are P is as little undermined by the *non*epistemic fact that the β's are non-P as it receives evidential support from the *non*epistemic fact that the α's are P. It is an essential feature of induction that it allows for epistemically legitimate inferences from *known* premises to conclusions that are *false* – if only their falsity is not within our ken.

[5] I focus on what John Pollock calls a *rebutting* defeater (Pollock 1984, 424; cf. his 1990 and 1994). Pollock's *undercutting* defeaters will not play a role in this essay. Nor will I here discuss what might be called 'probabilistic' defeat, generated by knowledge that the β's are only *probably* not-P.

question would thus still be projectible. Yet although predicates are only *conditionally* discriminating, some are *universally* discriminating, i.e., discriminating under all circumstances whatsoever. Consider my favorite example.[6] Enumerative induction is *defined* as projection from the members of I_α to those of I_β. Focus now on the predicate "sampled", which has a stipulated use: it applies to all and only the α's, i.e., to elements of I_α exclusively.[7] If we know the definition of enumerative induction, then once we know that we are to perform an inductive inference we know this predicate to apply to all and only the elements of I_α. As a result, "sampled" satisfies conditions (i) and (ii) for arbitrary induction sets and is therefore *universally* discriminating. The hypothesis that all *K*'s are sampled is always perfectly supported by our inductive evidence, while it is always defeated as well. The predicate "sampled" might therefore be called "unprojectible" *simpliciter*. However, this characterization should not blind us to the fact that, strictly speaking, unprojectibility is not an intrinsic feature of this predicate, but a consequence of the epistemic context it creates.

The puzzles to be discussed can be constructed around any discriminating predicate. To indicate their universality (and to tie in with Goodman's illustration of his paradox), I will predominantly base the puzzles on the predicate "sampled", which therefore plays a prominent role in this paper. Yet it is important to keep in mind that these puzzles could also be framed around a non-universally discriminating predicate such as "plastic", given an epistemic context in which this predicate is itself discriminating. Variation of the examples with respect to the discriminating predicate used does not affect the logical structure of the problem cases to be discussed.

2. The disjunctive riddle

While any plausible theory of confirmation easily accommodates discriminating predicates, whether universal or not, disjunctive weakening generates a problem. As shown in the introduction, disjunctive weakenings of discriminating predicates seem to permit drawing quite arbitrary conclusions from obviously irrelevant evidence. Even worse, they generate straightforward contradictions:

> CASE 1: From our knowledge that the 99 balls are sampled, we deduce that they are sampled or green and that they are sampled or non-green (e.g., sampled or blue). But of course we know that the final ball does not belong to the evidence class I_α; thus there is a defeater for the projection of "sampled".

If we were licensed to project both disjunctive hypotheses, we would be permitted to draw the disastrous conclusion that the last ball is, and isn't, green. But clearly we are *not* so permitted, on pain of contradiction, which indicates the failure of the *joint* projectibility of the disjunctive hypotheses.

The fact that the two disjunctive hypotheses are individually confirmed but not jointly projectible is no more mysterious than the unprojectibility of discriminating predicates themselves. It is fully explained by reference to background knowledge.

[6] The example is essentially due to Goodman (1946) and Leblanc (1963).

[7] Of course, "sampled" may be taken to refer to a *type* of property rather than a property itself, since I_α varies with the case. I refrain from discussing such subtleties here.

Indeed, the *conjunction* of these disjunctions, i.e., "sampled or green, and sampled or non-green" is logically equivalent to "sampled" and therefore itself (universally) discriminating. But merely uncovering the source for the failure of *joint* projectibility does not solve our problem. Reference to the defeater alone does not enable us to determine whether one of the disjunctive predicates remains projectible and, if so, which one enjoys this privilege: none of the disjunctive hypotheses conflicts individually with background knowledge.

This is not to say that we have any practical doubts whatsoever. Our intuitive verdict on what is and what is not projectible is perfectly clear: *None* of the conflicting hypotheses can be legitimately maintained in this situation. From the known fact that the α's, but not the β's, are sampled, we can neither infer that the β's are green, nor that they are blue. The two conflicting hypotheses are both unsupported in a substantial sense of the term. The problem then is not what to say about such cases, but how to justify the verdict. In contrast to the simple case discussed in section 1, the defeater generates the disjunctive riddle, yet it does not by itself put us in a position to solve it.

In reaction to this embarrassment we might be attracted to a simple solution, according to which we expand the scope of defeat and ban disjunctive weakenings together with the discriminating predicates themselves. The unprojectibility of, say, "sampled or green" is then an immediate consequence of the unprojectibility of "sampled" and the fact that the former is derivable from the latter by disjunctive weakening. To conceptually distinguish the two cases, I will speak of *direct* defeat (of "sampled") and of *derivative* defeat (of "sampled or green"), respectively. Derivative defeat, so conceived, would liberate us from any threat of a contradiction, indeed from all of the mentioned troubles. If neither "sampled or green" nor "sampled or non-green" were projectible, there would be no inference to the color of future objects and hence no riddle in the first place.

Simple and effective as this unrestricted conception of derivative defeat might be, it suffers from serious shortcomings. One worry is that it is utterly *ad hoc*, unjustified by any rationale other than our wish to avoid disastrous consequences. We are in need of some independent justification, some explanation of why a logically weaker predicate should be abandoned along with its stronger ancestor. In response to this difficulty, one might resort to non-epistemic criteria and deny projectibility to the disjunctive predicates under consideration on the grounds that they are, say, positional (Carnap 1947 and 1971), insufficiently entrenched (Goodman 1983), or non-natural (Quine 1969; Lewis 1983). But this maneuver would oblige us to explain, first, the respective characterization and, second, why the respective feature should be detrimental to inductive inference. I think it is fair to say that none of the explanations on offer are fully satisfactory in these respects.

There is an even more serious objection to an unrestricted conception of derivative defeat. If *all* disjunctive weakenings of discriminating predicates were themselves unprojectible, induction would be impossible. Consider any predicate P which you deem projectible ("green", "conducts electricity", "has half-integer spin") with respect to a given induction set. If an object is P, then it is P *or* Q for arbitrary Q's. "Green *or* sampled", for instance, constitutes a disjunctive weakening of "green", but also of the discriminating predicate "sampled". Given the plausible principle that a hypothesis is defeated if some known logical consequence is, however, "green" is defeated if "green or sampled" is.[8] So we cannot simply extend the scope of defeat to all such disjunctive weakenings without thereby undermining the very possibility of induction as such. The predicate "sampled or (non-)green" must be projectible, at least sometimes.

At this stage of the dialectics, our quandary has become manifest. The rather innocent case of discriminating predicates has developed into a serious dilemma via the simple and logically impeccable operation of disjunctive weakening. Our notion of induction generates contradictions – if we do not supplement direct defeat with a stronger criterion and allow for some form of derivative defeat. Or else it has an empty extension – if there is no restriction on derivative defeat. There remains only one way of escaping the dilemma, namely, to steer a middle course: sometimes disjunctive weakenings of discriminating predicates are projectible, sometimes they are not. Our problem then is to identify the precise conditions under which disjunctive weakenings of discriminating predicates are derivatively defeated.

To fix our intuitions, let me discuss a second scenario. With respect to CASE 1 we came to the conclusion that none of the two disjunctive hypotheses is projectible. We can predict neither that the final ball is sampled or green, nor that it is sampled or blue. Compare, however, CASE 2 which adds to CASE 1 that we have looked at the samples and found them green.

> CASE 2: We know the 99 balls to be sampled because we know that they are used as the basis for an inductive inference. From this we derive that they are sampled or green and that they are sampled or non-green (e.g., sampled or blue). We have also carefully examined the sample balls and observed that they are green. As before, we also know, of course, that the final ball is not sampled.

In contrast to CASE 1, we have not been empirically idle in CASE 2, which allows us to draw interesting conclusions about the final ball; we confidently, though of course only fallibly, predict its color. In fact we have overwhelming genuine confirmation for the hypothesis that the last ball is green. Our inductive inference can take either of two ways. We may project the predicate "green" and arrive directly, via simple enumerative induction, at the hypothesis that all balls are green. Alternatively, we may arrive at the very same conclusion by way of the following (admittedly circuitous) form of reasoning: "Because the samples are green, they are green or sampled. This supports the prediction that *all* balls from the urn, the final ball included, are green or sampled. But as the final ball is not in the evidence set, it must be green." The projection of "sampled or green", excluded in CASE 1, is genuinely

[8] Consider only the following argument for the mentioned principle. If the projection of a hypothesis H is defeated it cannot be maintained by S. Given that S knows that H^* entails H, S cannot accept H^*, either. Thus a defeater for the projection of H also defeats the stronger H^*, if only a suitable closure principle is presupposed. Of course we might be tempted, possibly inspired by the debate on the closure of knowledge (cf. Dretske 1971 and Nozick 1981), to deny that projectibility is closed under known disjunctive weakening and, therefore, that defeat does not transfer to logically stronger predicates. But this is an unpromising line to follow. Firstly, it will not do to deny the principle; one would have to deny all relevant instances of the principle, which is simply implausible. Secondly, even if projectibility is not closed under *known* disjunctive weakening, there are plausible modifications (cf., e.g., Williamson 2000, 117, and Freitag 2013, 87–91, for alternative versions of the closure principle for knowledge). Finally, it is highly desirable to formulate a solution for our problem which allows for closure. Whatever may go wrong in the disjunctive riddle and in Goodman's paradox, surely we should not blame the excessive use of *de*duction.

confirmed and hence perfectly legitimate in CASE 2. Our challenge is precisely to explain this fact.

3. The solution

Up to this point, we have considered as relevant only the *contents* of our (inductive and total) evidence. I will now suggest that it does not only matter *what* we believe, but also *why* we believe what we believe. The solution to our riddle will be based on considerations of *epistemic dependence*. Let me first introduce the notion by reference to another case, one that is independent of any problem of induction.

I believe that Kurt Gödel is a great logician or a great lover. I believe this for a single reason: I believe him to be a great logician. If I were to obtain proof of the fact that Schmidt, not Gödel, had discovered the incompleteness of arithmetic, I would not only abandon the belief that Gödel is a great logician, but also the disjunctive belief that came with it: I have no evidence concerning Gödel's amorous abilities.[9] But imagine that Gödel led a double life: mathematician during the day, Don Juan at night. And so there is Babette, who also believes Gödel to be a great logician or a great lover, although she is utterly oblivious about his scientific accomplishments. If she were told that the incompleteness theorem is due to the hapless Schmidt, she would *not* abandon her disjunctive belief about Gödel. So there are two persons, Babette and me, holding the very same (disjunctive) belief for very different reasons. *My* disjunctive belief epistemically depends on my opinion about his scientific achievements, while Babette's selfsame disjunctive belief depends on her evaluation of his erotic pastime. The difference in epistemic dependence perfectly explains why we react differently to the very same new information.[10]

Epistemic dependence plays a crucial role in what follows, yet its proper context is that of *deductive* inference (as witnessed by the Gödel example). I will therefore, and because our illustration provides a sufficient grasp of this notion, not aspire to a precise definition of the term here.[11] Let me emphasize, however, that relations of epistemic dependence can alter with the advent of new evidence and hence must not be mistaken as being determined by the method of original belief acquisition. In particular, epistemic dependence is not constituted by *actual* deductive inference. Suppose that I have obtained my beliefs about Gödel years ago during my logic

---

[9] Nor is my disjunctive belief itself 'basic', i.e., epistemically independent of the disjunct beliefs.

[10] Sven Ove Hansson (1999, 19) has discussed an example similar to the Gödel case in connection to AGM belief revision theory and provides essentially the same solution that I have been discussing in the main text: We have to distinguish between basic beliefs and dependent beliefs; belief update is then restricted to the set of basic beliefs.

[11] In the main text I provide a provisional characterization of epistemic dependence in terms of counterfactual dependence: my disjunctive belief about Gödel is epistemically based on my belief about his logical genius since I *would* discard it *were* I to find out that 'his' famous proof is actually by poor Schmidt. But this explanation must not be mistaken for an analysis. I am convinced that an approach in terms of counterfactuals gets us a long way in an attempt to capture or point to epistemic dependence, but I am not sure whether it can be turned into a fully satisfactory definition: too intricate are the ways of counterfactuals. Indeed I take it to be an open question whether epistemic dependence can be fully explained by reference to some other concept (e.g., counterfactual dependence, justification, the basing relation) at all. It might be a primitive notion. At any rate, my aim in this paper is not to give an account of epistemic dependence, but to point to its role in certain puzzles of induction. For these very limited purposes it suffices to have a pre-theoretical grasp of the application conditions of this notion with respect to the cases of interest.

education, but that recently I have come into possession of Babette's diary, which contains ample and convincing reference to Gödel's love life. If I were *then* informed that Schmidt, not Gödel, was the one who conceived of the famous proof, I would still be shocked, but my disjunctive belief would not be shaken. My newly acquired belief that Gödel is a great lover would render my belief that Gödel is a great logician or a great lover independent of its original basis, i.e., independent of the belief from which it has originally been deduced.

Back to the disjunctive riddle. CASES 1 and 2 differ with respect to the projectibility of but one disjunctive hypothesis, namely, that all balls are sampled or green. This difference can be straightforwardly traced to a difference in epistemic dependence. Our respective epistemic positions in CASES 1 and 2 are, prior to any inductive inference, comparable to mine with respect to Gödel – before and after I was enlightened by Babette's diary. In CASE 1, we have no opinion on the color of the sample balls, whence "green" figures in the complex "sampled or green" merely as an evidentially underdetermined disjunct. That is, in CASE 1 the disjunctive evidence epistemically depends on the knowledge that the samples are sampled: if we did not possess any such knowledge (e.g., because we did not intend to inductively infer from these samples at all), we would not believe that they are sampled-or-green either. In CASE 2, however, we have observed the colors of the samples and are hence in possession of another, independent, basis for the selfsame disjunctive evidence. That the samples are sampled or green can be derived from the observationally known fact that they are green. Our disjunctive evidence then is epistemically *independent* of the evidence for "sampled". If we then did not know the sample balls to be sampled, we would continue to believe that they are sampled or green, because we would continue to believe that they are green.

The difference in epistemic dependence would be wholly immaterial, were it not for the fact that "sampled" is a *discriminating* predicate and hence unprojectible itself:[12] evidence that is epistemically dependent on the evidence concerning a discriminating, and therefore unprojectible, predicate is itself rendered inductively inert, unable to genuinely confirm inductive hypotheses. This yields the following results. In CASES 1 and 2 the disjunctive evidence for the projection of "sampled or blue" is epistemically dependent on the evidence for the projection of the (discriminating) predicate "sampled", which renders the projection of "sampled or blue" not genuinely confirmed and therefore derivatively defeated in both scenarios. The projection of "sampled *or green*", on the other hand, is derivatively defeated only in CASE 1. In CASE 2 it is unaffected by the defeater, as the disjunctive evidence is then independent of that concerning "sampled". The pertinent difference between the projections of "sampled or green" in CASES 1 and 2 is due to a difference in epistemic dependence – made relevant by the presence of a defeater.

Extrapolating from this particular example, we can formulate sensible restrictions on derivative defeat on the basis of the notion of epistemic dependence: *An (inductively confirmed) hypothesis is derivatively defeated if and only if the pertinent inductive evidence epistemically depends on the inductive evidence for the projection of a discriminating*

[12] In the absence of a defeater we would in both cases be perfectly licensed to conclude that all balls are sampled or green. But then again, we would not be able to deduce from *this* hypothesis that the final ball is green! In CASE 1, thus modified, we would not be in a position to infer its color at all (because we would *not* know that the final ball is not sampled). And in CASE 2, equally modified, we would only be able to entertain the green-hypothesis by projecting "green" directly, i.e., without the intermediary step of disjunctive weakening.

predicate.[13] Derivative defeat thus constrained allows us to steer the required middle course. The induction-undermining power of the defeater extends to disjunctive weakenings of discriminating predicates only if the former are epistemically dependent on the latter. Armed with this result, we turn to the grue-paradox.

4. Grue

Define "grue" as follows. It applies to an object iff it is either sampled and green, or not sampled and blue.[14] Suppose now that we have examined the 99 sample balls and found them green (i.e., we are in an evidential situation just like the one described in CASE 2). Our evidence then yields perfect inductive support for the general hypothesis that *all* balls from the urn are green. Moreover, if the samples are green, they are by definition also grue. That is, our evidence equally confirms the hypothesis that all balls, the final one included, are grue. Yet obviously we cannot maintain *both* hypotheses, since an unsampled grue ball is blue, not green. This constitutes Goodman's paradox.[15]

As in CASES 1 and 2, the competing hypotheses are rendered jointly unprojectible only because of undermining background knowledge, and the defeater does not, at least not by itself, determine whether either of the conflicting hypotheses is projectible and, if so, which one enjoys this privilege. Moreover, the conflicting predictions are equally supported by the inductive evidence: we know that the samples are green and that they are grue. There is no difference between the respective types of evidential support when it comes to the *content* of the evidence alone. Yet, like in CASE 2, and in contrast to CASE 1, there is a difference with respect to *epistemic dependence* between the different kinds of evidential beliefs. We trivially know that the sample balls are sampled. And we have *seen* that they are green, because of which we are in possession of independent knowledge about their color. But that the samples are grue has only been *inferred* from the knowledge that they are green *and* that they are sampled: if we did not know that the 99 balls are sampled, we would also fail to believe that they are grue. The inductive evidence for the grue-hypothesis therefore epistemically depends on the evidence for a defeated hypothesis. The grue-hypothesis is derivatively defeated according to our criterion. The competing green-hypothesis, on the other hand, is legitimately projectible due to the constraints on derivative defeat formulated above. In the evidential situation described by Goodman, we base our evidence that the samples are green on visual inspection, so it would not be lost if we did not believe that the sample balls are sampled. Green-evidence genuinely confirms the green-hypothesis.

A slightly more elaborate analysis of the grue-paradox reveals its close connection to the disjunctive riddle discussed above. "Grue" is defined as "green and sampled, or non-sampled and blue", which is analytically equivalent to "sampled or blue, and green or non-sampled". That is, "grue" can be construed as a conjunction of two

[13] Observe that, given this definition and the assumption that epistemic dependence is reflexive, direct defeat is but a special case of derivative defeat.

[14] In contrast to Goodman 1983, I define "grue" in terms of "sampled" instead of "examined before *t*". Observe also that, if we substitute "non-green" for "blue", "grue" is definable as "green if and only if sampled".

[15] My description in the main text is in all relevant respects faithful to the famous passage from Goodman 1983, 73–74.

disjunctions, one of which, the predicate "sampled or blue", is the disjunctive weakening of a discriminating predicate. Let's focus on *this* disjunctive predicate. In Goodman's scenario, as in Cases 1 and 2, we know that the samples are sampled or blue only because we know them to be sampled. The evidence for this disjunctive projection epistemically depends on that of a discriminating predicate and is therefore derivatively defeated. This result is directly relevant for the grue-case. Plausibly, if we can't predict that all balls are sampled or blue, we can't project the logically stronger "sampled or blue, and non-sampled or green" either, which defines "grue".[16] The projection of "grue" is derivatively defeated for the same reason that undermines the projectibility of "sampled or blue". The projection of "sampled or green", however, is epistemically *independent* of the discriminating predicate and hence unaffected by the presence of the defeater. "Sampled or green" remains projectible, and so is "green" itself. This explains and justifies our favoring "green" over "grue". And it shows that, at bottom, the grue-problem is but a version of the disjunctive riddle.[17]

"Green" prevails over "grue" because the projection of the latter, and only the latter, is derivatively defeated. If this is correct, there is no need for a general distinction between projectible and unprojectible predicates. At the very least Goodman's riddle does not call for such a distinction. Goodman's paradox turns out to be a purely epistemic problem, generated and resolved by epistemic context – including relations of epistemic dependence – alone.

To strengthen this result – and to further clarify it – let me address some possible worries and objections. *First objection*: If the unprojectibility of "grue" is contingent upon (1) the presence of a defeater, and (2) the fact that grue-evidence epistemically

[16] I here assume again that a predicate is projectible only if its (known) logical consequences are too (cf. also fn. 8). But note also that, if a belief that p is epistemically dependent on any of its (known) logical weakenings, the mentioned assumption is superfluous.

[17] Let me only in passing compare my own solution to Frank Jackson's, which is based on the following "counterfactual condition" on projectibility: "Certain K's which are F being G does not support other K's which are not F being G if it is known that the K's in the evidence class would not have been G if they had not been F" (Jackson 1975, 123; variables changed). We know, Jackson says, that the sample balls would not have been grue (G) if they had not been sampled (F), since we know that they would have been green, but not blue. So we are aware of the fact that the grueness of the sample balls, but not their greenness, depends on their being sampled, which, according to Jackson, explains our preference for the green-hypothesis over its competitor in cases of conflict, e.g., if there is a further, unsampled ball. Jackson's account looks promising with respect to "grue" and it appears to incorporate precisely those features which I have stressed in my own solution: (direct) defeat and some sort of (counterfactual) dependence. Yet there is a crucial difference, which – ignoring all subtleties – can be presented as follows: Jackson's solution rests on (known or believed) dependence between facts (described by the following scheme: $\mathrm{B}(p \mathrel{\Box\!\!\rightarrow} q)$), while mine is based on dependence between beliefs ($\mathrm{B}p \mathrel{\Box\!\!\rightarrow} \mathrm{B}q$). That Jackson's proposal fails can best be brought out with respect to the disjunctive riddle. Recall that, in Case 1, we don't know anything about the color of the sample balls. Our only evidence is that they are sampled – and whatever we can derive from that fact, e.g., that they are sampled or green and that they are sampled or non-green. Surely, therefore, in this case the disjunctive hypothesis does not meet Jackson's criterion: we do *not* know that, if the samples had not been sampled, they would not have been sampled or green, because we don't know that they would not have been green. Nor do we know that they would not have been (sampled or) blue. Since Jackson's condition is not satisfied for either of the two conflicting predicates, their individual projections are not ruled out and the problem remains. Jackson's theory is of very limited value at best, unable to solve the disjunctive riddle, a central and important puzzle about projectibility. Moreover, if the grue-puzzle is but another version of the disjunctive riddle, Jackson's account does not provide the right solution to Goodman's paradox either. For essentially the same reason, I object to the modified version of Jackson's position in Jackson and Pargetter 1980.

depends on sampled-evidence, it seems that there are situations in which grue-style predicates are projectible. But this is contrary to sound intuition. "Grue" can never be projected. *Reply:* I claim indeed that *if* (1) or (2), or both, do not obtain, there is no reason to abandon the grue-hypothesis. In the absence of a defeater, there is no epistemic conflict between "green" and "grue", and so both are, other things being equal, projectible. And if there is a conflict, but instead of the grue-evidence it is the green-evidence that epistemically depends on the sampled-evidence, then the green-hypothesis is to be discarded and the grue-style hypothesis prevails. But I am only committed to this conditional claim; I am not committed to the claim that it is possible that (1) and (2) do not obtain.

But let me speculate. Unlike "plastic", "sampled" is *universally* discriminating, so the grue-paradox is universal itself: with respect to any induction set, there is a conflict between the green- and the grue-hypothesis.[18] So "grue" cannot be projectible on the grounds that there is no conflict with "green" at all. Whether there may be cases in which grue-evidence is not dependent on sampled-evidence is a more interesting question, to which I am inclined to give a negative answer. At least, I have not been able to construct any plausible example in which green-evidence, but not grue-evidence, is epistemically dependent on sampled-evidence. We typically know the samples to be sampled by reflecting on the fact that we are about to perform an inductive inference. We also typically investigate objects for their weights and heights, for their material constitutions and colors. Such knowledge is usually acquired by observation, testimony, etc., this then forming another part of our (comparatively) 'basic' evidence. Grue-evidence, on the other hand, is then merely derived from the evidence that the samples are green and that they are sampled. That the unprojectibility of "grue" is doubly conditional does hence not entail the existence of realistic scenarios in which these conditions are not satisfied. If, however, contrary to my present suggestion, grue-favoring evidence scenarios should exist, they are of a fairly strange sort – *very* different from the situations in which we usually find ourselves and which inform our pre-theoretical intuitions. This answers the first objection.

Second objection. "Grue" is defined in terms of "green", "blue", and "sampled". But, as Goodman (1983, 79–80) has so convincingly argued, we can also imagine or construct other languages, in which "grue" and "bleen" (defined in English as "sampled and blue, or non-sampled and green") are primitive, while "green" in turn is defined in terms of these predicates: an object is green iff it is sampled and grue, or unsampled and bleen. That is, the objection goes, we can imagine situations in which "grue" is conceptually prior to "green". Yet unless it is a language-relative feature, projectibility must not depend on questions of conceptual priority. *Reply:* Firstly, a clarification. Our definition of "grue" has the function of introducing a new predicate into ordinary or philosophical English. Nothing else is intended or implied. In particular, the definition does not have the purpose of indicating or establishing relations of relative conceptual priority or posteriority. Indeed, the semantical relations between the predicates described in this paper do not, in general, indicate relations of conceptual priority or posteriority at all: even in languages in which "grue" is conceptually prior to "green" etc., it is logically equivalent to "green and sampled, or non-sampled and blue". (Conversely, the predicate "green" is logically equivalent to "sampled and grue, or unsampled and bleen" in English.) Secondly, and

[18] Again, I here ignore pathological cases in which we don't know that the future objects are not sampled.

more importantly, our criterion for derivative defeat does not refer to the notion of conceptual priority. The sole arbiter for derivative defeat is epistemic dependence on evidence concerning discriminating predicates. Epistemic dependence, however, is logically unrelated to conceptual dependence. Even if “green” were *conceptually* posterior to “grue”, its *epistemic* priority (of the sort required) in evidential situations such as the one described by CASE 2 would not be affected: we would still know that the samples are grue only because we know that they are green and sampled.[19] Questions of conceptual dependence are therefore irrelevant for issues concerning epistemic dependence and hence for questions of derivative defeat.

According to the *third objection* it is doubtful whether our criterion for derivative defeat can distinguish between “green” and “grue” at all. It claims that if the grue-hypothesis is derivatively defeated then the green-hypothesis is too. The reasoning, I imagine, would run as follows. According to the criterion proposed, a projection is derivatively defeated if the pertinent evidential beliefs epistemically depend on the evidence for a (directly) defeated projection. This eliminates the grue-hypothesis. But, the argument continues, this also rules out the green-hypothesis: if the 99 balls had not been sampled, we wouldn’t be in possession of *any* inductive evidence concerning them, let alone evidence that the sample balls are green. As a consequence, it will be argued, the projection of “green” depends, like that of “grue”, on the fact that these balls are sampled. It is therefore derivatively defeated as well. Instead of establishing an asymmetry between the two conflicting hypotheses, the proposed solution in terms of derivative defeat bans both of them and ultimately leads to the unpalatable consequence that no predicate is ever projectible.

Reply: First, it should be noted that objects may have been examined without their *being* sampled in my definition of the term: To be sampled means to be part of the inductive basis I_α for an (enumerative) induction, but obviously we can acquire knowledge about objects without using them as a basis for inductive inference. But even if, for the sake of the argument, we were to concede that we would not know that the 99 balls are green if we had not sampled them, the objection would be ineffective. To see this, suppose that all of our evidence originates in the mere *fact* that these balls were sampled, and that we would not have this evidence had we not sampled these balls. No doubt this would constitute a form of *genetic* dependence of our evidence on the *fact* that they have been sampled, but it would not constitute *epistemic* dependence: we would continue to believe the samples to be green even if we did not *believe* that 99 sample balls are in fact sampled. Yet epistemic *independence* (from the discriminating predicate) secures *ceteris paribus* the projectibility of the green-hypothesis. Objection 3 is based on a misunderstanding or misapplication of the criterion for derivative defeat, or more specifically, on a confusion of genetic and epistemic dependence.

5. Conclusion

Let me conclude with a reflection on what may have been the major impediment to a proper construal of Goodman’s paradox. In Goodman’s original examples, the paradox is framed around a *universally* discriminating predicate (“sampled” or, in Goodman 1983, “examined before *t*”). This has the advantage of demonstrating the

[19] It is an entirely different matter whether evidence scenarios like that described by Goodman are possible in situations in which “grue” is conceptually prior to “green”.

generality of the problem, but also the drawback of disguising the role of discriminating predicates in the construction of the problem: it is trivial and hence often goes without saying that the α's are distinguished from the β's with respect to this very property. That "sampled" has a function first in confirming the grue-hypothesis and second in defeating it is therefore more often overlooked or ignored than emphasized. Neglect of the role of the discriminating predicate is not a minor lacuna only. It precludes the very possibility of a proper solution by obscuring the role of the defeater and hence by disguising universal unprojectibility as unconditional, intrinsic unprojectibility. And it prevents recognition of the role of epistemic dependence. With his ingenious choice of example Goodman has unwittingly concealed those elements of the epistemic background that are crucial to the case.

The purpose of this paper was to bring the epistemic background into the theoretical foreground and to show that Goodman's paradox is generated by the presence of a defeater, and solvable by reference to epistemic dependence. Once we realize that epistemic context plays an essential role and that it comprises not only our total evidence, but also the epistemic relations between evidential beliefs, the solution to the paradox is surprisingly simple. The projection of "grue" is, and that of "green" is not, derivatively defeated, because grue-evidence is, while green-evidence is not, epistemically dependent on the 'evidence' that the samples are sampled. In the light of epistemic context, therefore, Goodman's paradox loses much of its gruesome appearance.[20]

[20] I am grateful to three anonymous referees of this journal and to the members of the group that must not be named, especially Moritz Schultz (who must be named), for a number of helpful suggestions. Christopher von Bülow proof-read the whole piece. My greatest thanks go to Wolfgang Spohn and Alexandra Zinke, who patiently accompanied the development of the notion of derivative defeat.

References

CARNAP, Rudolf, 1947, "On the Application of Inductive Logic", *Philosophy and Phenomenological Research* 8, pp. 133–147

—, 1971, "A Basic System of Inductive Logic", in: R. Carnap and R. C. Jeffrey, eds., *Studies in Inductive Logic and Probability*, Vol. I, Berkeley CA and Los Angeles: University of California Press, pp. 33–165

DRETSKE, Fred, 1971, "Conclusive Reasons", *Australasian Journal of Philosophy* 49, pp. 1–22

FREITAG, Wolfgang, 2013, *I Know. Modal Epistemology and Scepticism*, Münster: mentis

—, 2015. "I bet you'll solve Goodman's Riddle", *Philosophical Quarterly* 65, pp. 254–267

GOODMAN, Nelson, 1946, "A Query on Confirmation", *Journal of Philosophy* 43, pp. 383–385.

HANSSON, Sven Ove, 1999, *A Textbook of Belief Dynamics: Theory Change and Database Updating*, Dordrecht: Kluwer Academic Publishers

JACKSON, Frank, 1975, "Grue", Journal of Philosophy 72, pp. 113–131

—, and PARGETTER, Robert, 1980, "Confirmation and the Nomological", *Canadian Journal of Philosophy* 10, pp. 415–428

LEBLANC, Hugues, 1963, "A Revised Version of Goodman's Confirmation Paradox", *Philosophical Studies* 14, pp. 49–51

LEWIS, David, 1983, "New Work for a Theory of Universals", *Australasian Journal of Philosophy* 61, pp. 343–377

NOZICK, Robert, 1981, *Philosophical Explanations*, Oxford: Oxford University Press

POLLOCK, John L., 1984, "A Solution to the Problem of Induction", *Noûs* 18, pp. 423–461

—, 1990, *Nomic Probability and the Foundations of Induction*, Oxford and New York: Oxford University Press

—, 1994, "The Projectibility Constraint", in: Stalker 1994, pp. 135–151

QUINE, W. V. O., 1969, "Natural Kinds", in his *Ontological Relativity and Other Essays*, New York: Columbia University Press, pp. 114–138

STALKER, Douglas (ed.), 1994, *Grue! The New Riddle of Induction*, Chicago and La Salle IL: Open Court

WILLIAMSON, Timothy, 2000, *Knowledge and Its Limits*, Oxford: Oxford University Press