

On the Justification of Deduction and Induction

Franz Huber

Department of Philosophy

Institute for the History and Philosophy of Science and Technology

University of Toronto

Penultimate version: please cite the paper that is forthcoming in the
European Journal for the Philosophy of Science

Contents

1	Introduction	3
2	Hume's Argument against Induction	4
3	Carnap's Rationalist View of Rationality	7
4	The Continuum of Carnap's Rationalism	13
5	Goodman's Fiction about the Facts of Forecast	16
6	Reichenbach Rules	20
7	Haack's Argument against Deduction	27

Abstract

The thesis of this paper is that we can justify induction deductively relative to one end, and deduction inductively relative to a different end. I will begin by presenting a contemporary variant of Hume (1739; 1748)'s argument for the thesis that we cannot justify the principle of induction. Then I will criticize the responses the resulting problem of induction has received by Carnap (1963; 1968) and Goodman (1954), as well as praise Reichenbach (1938; 1940)'s approach.

Some of these authors compare induction to deduction. Haack (1976) compares deduction to induction, and I will critically discuss her argument for the thesis that we cannot justify the principles of deduction next. In concluding I will defend the thesis that we can justify induction deductively relative to one end, and deduction inductively relative to a different end, and that we can do so in a non-circular way. Along the way I will show how we can understand deductive and inductive logic as normative theories, and I will briefly sketch an argument to the effect that there are only hypothetical, but no categorical imperatives.

1 Introduction

The thesis of this paper is that we can justify induction deductively relative to one end, and deduction inductively relative to a different end. Here is how I will arrive at this unorthodox conclusion. I will first present a contemporary variant of Hume (1739; 1748)'s argument for the thesis that we cannot justify the principle of induction. In doing so I will stress that the resulting problem of induction is problematic only insofar as we have cognitive limitations and ends. (I assume ultimate cognitive ends to be intrinsic cognitive desires, but nothing in this paper hinges on this assumption.) We will not solve the problem of induction if we do not pay due respect to our cognitive limitations and ends, as it is only because of the latter that the problem of induction is a problem in the first place.

I will then criticize the responses the problem of induction has received by Carnap (1963; 1968) and Goodman (1954). Logic, whether it is deductive or inductive logic, can be understood as a normative theory. The upshot of the discussion of Carnap's work will be that we get the justification of our normative theories wrong if we have the wrong theory of normativity. The discussion of Goodman's work will drive home the point that we get the inductive case wrong if we get the deductive case wrong. This critical discussion of the literature will turn into a constructive proposal once we have incorporated Reichenbach (1938; 1940)'s insight that the justification of norms or imperatives is relative. Logic, like every (set of) norm(s) or imperative(s), is hypothetical in nature: its validity consists in being the means to attaining some cognitive end. Which logic we ought to follow depends on what our cognitive limitations and ends are. If we have no cognitive ends to begin with, we need not obey any logic, and the problem of the justification of logic – whether it is inductive or deductive logic – vanishes. If we have cognitive ends, then we need to make these ends explicit before we can even meaningfully ask whether a given logic is justified *relative to* these ends.

Carnap and Goodman compare induction to deduction. Haack (1976) compares deduction to induction. I will critically discuss her argument for the thesis that we cannot justify the principles of deduction next. I will conclude by defending the thesis that induction is justified by a deductively valid argument relative to one cognitive end, and deduction is justified by an inductively strong argument relative to a different cognitive end. Since the two cognitive ends relative to which we justify induction and deduction, respectively, are different, the justification of each of these two logical systems is non-circular. It is not claimed that the conjunction of deduction and induction is justified (by a deductively valid, inductively strong, or any other argument) relative to the conjunction of these two ends.

2 Hume's Argument against Induction

Here is a contemporary variant of Hume's (1739; 1748) argument for the thesis that we cannot justify the principle of induction, whatever its precise form.

Premise 1: we can justify the principle of induction only if there is a deductively valid or inductively strong argument which does not presuppose its conclusion, whose premises are restricted to information we have, and whose conclusion says that the principle of induction usually leads from true premises to true conclusions.

Premise 2: there is no deductively valid argument which does not presuppose its conclusion, whose premises are restricted to information we have, and whose conclusion says that the principle of induction usually leads from true premises to true conclusions.

Premise 3: there is no inductively strong argument which does not presuppose its conclusion, whose premises are restricted to information we have, and whose conclusion says that the principle of induction usually leads from true premises to true conclusions.

Conclusion: we cannot justify the principle of induction.

This argument is logically valid. Let us briefly discuss its premises, and the ingredients of the latter. The principle of induction can be characterized as a method for answering questions to which one does not have the answer in light of information that one has. While in general overly simplistic, for the purposes of this paper the following principle of universal induction will do: from the premise that all objects about which one has enough information – i.e. whether they are F and whether they are G – are G if they are F , one may and ought to infer the conclusion that all F s are G . Here ' F ' and ' G ' are arbitrary predicates. An application of the principle of induction is an inductively strong argument. The numbers alluded to in the notion of inductive strength will not play a role in this paper.

Different philosophers will have different views on what precisely is information we have. Some might say the information we have are the sentences or propositions we know, or perhaps assume, to be true. Others might say the information we have are the sentences, or perhaps the true sentences, we believe to be true. As long as there is some information that we do not have, i.e. some question to which we do not have the answer, it does not matter much which of these options or their combinations, if any, we choose (see, however, the caveat below).

The conclusion of the sub-argument supposedly justifying induction says that the principle of induction usually leads from true premises to true conclusions. If there are infinitely many instances of the principle of induction, one needs to make some assumption (such as defining a measure over the infinitely many instances) to make sense of the claim that the principle of induction *usually* leads from true premises to true conclusions. I will suppose that this assumption is such that the information we presently have is not sufficient to determine already now that the principle of induction usually leads from true premises to true conclusions, even if all its future instances with true premises have a false conclusion.

Premise 3 is true, or so I will presume. It is so because any inductively strong argument for the conclusion that the principle of induction usually leads from true premises to true conclusions presumably has to be inductively strong in the sense of this very principle. Hence any such argument presupposes that the principle of induction usually leads from true premises to true conclusions. Yet this is precisely the conclusion we want to derive. So any inductively strong argument for the conclusion that the principle of induction usually leads from true premises to true conclusions presupposes its conclusion: it is circular. Therefore there is no inductively strong argument whose conclusion says that the principle of induction usually leads from true premises to true conclusions and which does not presuppose its conclusion (and hence no such argument whose premises are restricted to information we have).

The presumption that the inductively strong argument for the principle of induction has to be inductively strong in the sense of this very principle, and with it premise 3, need not be true if there are several principles of induction that can inductively justify each other (in each case with respect to the conclusion that the principle of induction under consideration usually leads from true premises to true conclusions): the first justifies the second which justifies the third which justifies the first. If there are only finitely many principles of induction, we are still caught in a circle. If there are infinitely many principles of induction, we need not be caught in circle. In this case other arguments are needed to establish the third premise. It is perhaps worth pointing out that the mere fact that we rely on infinitely many principles of induction in itself is not reason enough to despair. However, the burden of proof is with the defender of the inductive justification of induction, and the infinitist options that come to mind quickly run into problems (see Skyrms 2000: ch. 3).

Premise 2 is true as long as there is some information that we do not have, and as long as the assumption made to give precise meaning to the term ‘usually’ does not violate the above supposition. This is so because every question to which we

do not have the answer is, by definition, a question whose answer requires information which we do not have. The principle of induction is a *general* principle that applies to all questions, including those to which we do not have the answer. And it is a truth of logic that there is no deductively valid argument whose premises are restricted to information we have, and whose conclusion says something that is the answer to a question to which we do not have the answer. Therefore, as long as there is some information that we do not have, and the supposition mentioned above is met, there is no deductively valid argument whose premises are restricted to information we have, and whose conclusion says that the principle of induction usually leads from true premises to true conclusions (and hence no such argument which does not presuppose its conclusion).

Two points are worth noting. First, the problem of induction does not arise for omniscient beings who have as information all truths. If it is indeed the case that the principle of induction usually leads from true premises to true conclusions, then these beings have this information. And if it is *not* the case that the principle of induction usually leads from true premises to true conclusions, then these beings have *this* information. Either way there is a deductively valid argument which does not presuppose its conclusion, whose premises are restricted to information these beings have, and whose conclusion says that the principle of induction usually leads from true premises to true conclusions – or does not do so, as the case may be. *We* who are not omniscient cannot provide such a deductively valid argument, because we do not have enough information (here the choice of what we take to be information that we have matters somewhat). The problem of induction is *our* problem, not that of omniscient beings.

Second, the problem of induction vanishes for cognitively contented beings who have no curiosity or other cognitive ends. The problem of induction vanishes for such beings in the sense that, while perhaps still a problem, it is one that need not, and does not, bother such contented beings. For *us* who are not (cognitively) contented in this way the problem of induction is real and does not vanish, because we generally want to have the information that we do not have – and often we need to in order to survive. Curiosity may have killed the cat. However, the cat would have died much sooner if it had not been curious.¹

¹To be sure, we do not need to have information about how induction works, or if it can be justified, in order to apply induction and be good at doing so. However, we do need to have such information if we want to delegate the task of making inductive inferences to computers. Besides our philosophical curiosity I take this engineering task to be the major reason for the interest in the justification of induction. To the extent that philosophy is a normative discipline it is a branch of engineering (cf. Quine 1986).

Premise 1 is false. However, the premise 1 is not false because it makes strong internalist assumptions about justification. I happily grant these. Rather, premise 1 is false, because we can justify the principle of induction by a deductively valid argument which does not presuppose its conclusion, whose premises are restricted to information we have, but whose conclusion says something other than that the principle of induction usually leads from true premises to true conclusions. We will see in section 6 what this other feature is relative to which we can justify the principle of universal induction by a deductively valid argument.

3 Carnap's Rationalist View of Rationality

Carnap (1934; 1936; 1937; 1947) famously denied Kant (1902)'s dictum (from *The Critique of Pure Reason*) that there are synthetic truths that we can know a priori. Afterwards Carnap worked extensively on probability and inductive logic (Carnap 1950; 1952; 1962). However, with the possible exception of Carnap (1962: §41) it is only in response to Nagel (1963: 820f) and, especially, Kemeny (1963) and Burks (1963) that Carnap (1963b: 978) briefly discusses “the controversial problem of the *justification of induction*.”

For Carnap the principle of induction takes the form of an inductive logic. Hence for him the justification of induction boils down to justifying the axioms of his inductive logic. These axioms specify a confirmation function (in Carnap 1950 it is one single confirmation function; in Carnap 1952 it is an entire family of confirmation functions). These axioms, and their logical consequences, are analytically true, as Carnap stresses frequently (for instance, Carnap 1963a: 71f).

According to Carnap (1963b: 987) we should “understand the problem of the justification of induction as the question as to what kinds of reasons can be given for accepting the axioms of inductive logic.” These reasons take the form of “conditions of adequacy” (Carnap 1963b: 977). They “are based upon our intuitive judgments concerning inductive validity, i.e., concerning inductive rationality of practical decisions (e.g., about bets)” (Carnap 1963b: 978). From this he infers that “[i]t is impossible to give a purely deductive justification of induction,” and that these “reasons are a priori.”

The notion of analyticity defended by Carnap (1934; 1936; 1937) defines analytic truth syntactically as derivability in a logical system. Carnap (1947) defines analytic truth semantically as validity in a family of models. This difference does not matter for us (unless the “reasons” mentioned above are semantically valid, but not derivable syntactically, and there is no indication that this is so).

What does matter for us is that Carnap's claim that "[i]t is impossible to give a purely deductive justification of induction" implies, on his very own account of analyticity, that the "reasons" which we cite in support of the various axioms of inductive logic, and which are "based upon our intuitive judgments about inductive validity," are synthetic. It is important to note that this is not true of what these reasons are reasons *for*, namely the analytically true axioms of inductive logic. Since Carnap (1963b) also claims that these "reasons are a priori," it follows that the reasons, or conditions of adequacy, that justify us in accepting the axioms of inductive logic are synthetic truths that we can know a priori. The source of our a priori knowledge of these synthetic reasons, or conditions of adequacy, is "intuitive judgment."

This is a curious thing to say, especially for Carnap, who, in Carnap (1934; 1936; 1937; 1947), held that there are no synthetic truths that we can know a priori. Let us have a closer look to see what has gone wrong. Carnap follows Kemeny in regarding the following condition of adequacy to be "intuitively valid" (Carnap 1963b: 977): the confirmation function "*c* must define a system of bets that is [strictly] fair." (Kemeny 1963: 720) This condition of adequacy is the reason why we ought to accept the axioms of probability theory as axioms of inductive logic. The axioms of probability theory themselves are analytically true sentences of the object language (and presumably we can know them a priori). In contrast to this the condition of adequacy is a synthetic sentence of the metalanguage. Yet we can still know it a priori, because "questions of rationality are purely a priori" (Carnap 1963b: 983). In my opinion it is here, in his view of rationality, where things go wrong. Before explaining why I think so, let me briefly discuss why Carnap might have thought that "questions of rationality are purely a priori". He might have thought so because of another dictum of Kant (1902)'s (this time from *The Groundwork of the Metaphysics of Morals*), namely the idea that there are absolute norms or categorical imperatives.

On the view of rationality, or normativity, adopted in this paper there are no categorical imperatives. Needless to say, there are utterances of normative, i.e. prescriptive or evaluative, claims whose grammatical form is categorical. Indeed, I will give an example of such a grammatically categorical imperative momentarily. What the present view denies is that there is any categorical deontic reality that is referred to by such grammatically categorical imperatives. On the present view of normativity there are only means to ends. Means-end relationships include being a contributing cause and other potentially non-deterministic causal relationships that cannot be established by formal methods alone, as well as being logically equivalent and other logical relationships that can be so established.

Which ultimate ends we have is a factual matter. Which actions, or intentions to act, are means to which ends often is so as well. What is never a factual matter on the present view is which ultimate ends one ought to have. In what follows I will briefly sketch an argument for this view.

Some imperatives are hypothetical imperatives. Say that Emmy should (intend or will to) not smoke. The grammatical form of this imperative is categorical. It is valid, because not smoking is a means to attaining the end of not developing lung cancer, an end Emmy is hypothetically assumed to have. Now enter Occam's razor, which is itself a norm or imperative. Occam's razor urges us to not multiply entities without necessity. In particular, it requires us to avoid additionally postulating the existence of a second kind of imperative, the category of categorical imperatives – unless we have to. We do not have to, though. Even if there were imperatives which we did not yet understand as hypothetical imperatives, it would not follow that we would have to postulate a second category of categorical imperatives. Instead we could always try harder. Occam's razor in turn is not a categorical imperative, but a means to attaining the truth *efficiently* (Kelly 2007). And it just so happens that I desire to hold true and informative beliefs, and to arrive at them efficiently.

Of course, things would be radically different if there was an argument to the effect that there are categorical imperatives. That, however, is not even possible, for we will see that the principles of logic – *any* logic, that is – are hypothetical imperatives, and no argument can do without some logical principle or other.

This is why I do not believe in categorical imperatives. To the extent that you also aim at holding true and informative beliefs, and to arrive at them efficiently, you ought to do so as well. Not because I tell you so, but because doing so is a means to attaining your ends. (This is not to say that holding false beliefs cannot be a means to attaining ends other than holding true beliefs.)

Korsgaard (2008) argues that there can be no requirement to take the means to our ends unless there are also required ends. Let me briefly point out why her arguments do not undermine the above claim. First, all that Korsgaard (2008) shows is that there can be no categorical requirement to take some action unless there are also ends that are categorically required. She does not show that there can be no hypothetical requirement that is conditional on some end that itself is not required. Someone who denies the existence of categorical imperatives will expect and embrace this conclusion rather than conclude that there are ends that are categorically required: that there are only hypothetical imperatives, and no action that is categorically required, is what the denier of categorical imperatives – Korsgaard (2008)'s "empiricist" – has claimed all along!

Second, contrary to what Korsgaard (2008) claims the empiricist cannot only explain why the instrumental principle that one ought to take the means to one's ends can *motivate* us. That is done by the ends or desires a hypothetical imperative is conditional upon. She can also explain why it can *guide* us. In particular, we can violate the instrumental principle in two ways: by having false or incomplete beliefs about what the means to our ends are, and about the consequences of taking these means; as well as by failing to intend or will to take the means to our ends. What ends we have is not always transparent to us. Sometimes this is so, because we make mistakes and have false beliefs about the consequences of realizing, i.e. taking the means to attaining, our ends. Sometimes this is so, because we simply cannot foresee all the consequences of realizing our ends. This is true for ends that are means to other ends. It is also true for ultimate ends (if we can have more than one), as we may miss that the only way of realizing one ultimate end prevents us from realizing another ultimate end that is at least as important to us.

Korsgaard (2008: 10f) thinks the empiricist must hold that there are no, and cannot be any, violations of the instrumental principle. She considers an agent whose ends are "shaped by [...] terror, idleness, shyness or depression" (Korsgaard 2008: 12) and thinks the empiricist cannot hold that this agent acts irrationally. However, the beliefs (about the pains and pleasures that are actually provided by taking the means to her ends) that this agent acts upon are simply false. It is precisely because this agent could or even does have other, true beliefs, yet fails to act upon the latter, that we consider her to act irrationally when she lets depression stop her from working (she acts as if she believed the work was difficult when she could or even does have the other, true belief that it is not), shyness stop her from calling (she acts as if she believed calling was embarrassing when she could or even does have the other, true belief that it is not), idleness stop her from going downtown (she acts as if she believed going was inconvenient when she could or even does have the other, true belief that it is not), and terror stop her from riding the roller coaster (she acts as if she believed riding was dangerous when she could or even does have the other, true belief that it is not).

To be sure, this is different for *omniscient* agents who *should* take precisely these actions that they *intend or will to* take, unless they suffer from weakness of will. Furthermore, *omnipotent* agents *intend or will to* take precisely these actions that they *do* take. So agents who are omniscient and omnipotent and who do not suffer from weakness of will can, indeed, *not* violate the instrumental principle. Nor can such agents be guided by the instrumental principle. However, this is as it should be: omniscient and omnipotent agents who do not suffer from weakness of will do exactly what they should do. They have no need for any guidance.

Finally, the *Ought* in the instrumental principle that one *ought* to take the means to one's ends is not *derived* (Korsgaard 2008: 6f) from the *Is* of having some end. Rather, the *Ought* is *conditional upon* the *Is* of having some end. Hypothetical imperatives are not categorical imperatives that some hypothetical, or conditional, be obeyed. Nor are they hypotheticals, or conditionals, whose consequent ought to be obeyed if its antecedent obtains (or is obeyed). Hypothetical imperatives are a *sui generis* category, *somewhat* similar to conditional beliefs which also cannot be reduced to beliefs that some conditional is true, or some conditional whose consequent is believed if its antecedent obtains (or is believed).

Hypothetical imperatives or norms are great. For instance, they allow us to explain why norms have normative force. They have such force, because ultimately they are conditional upon ultimate ends – that is, subjective intrinsic desires – and desires have force, ranging from mere wants to genuine needs. In the purely theoretical domain of epistemology, the strength, or force, of our cognitive desires explains the normative force of our doxastic or cognitive norms. In this domain taking the means to our cognitive ends can come in conflict only with taking the means to other cognitive ends of ours, as we never leave the internal world that is our mind. Outside the purely theoretical arena, where your taking the means to your ends may conflict with my taking the means to my ends, things become more difficult, but not different: there are no required means without some desired ends. In contrast to this, we presumably need to postulate the existence of a special authority or deity, or of moral facts or real virtues or objective intrinsic values, or other ontologically dubious categories in order to explain how categorical imperatives or norms can have normative force.

However, which hypothetical imperatives hold, and which desires they are ultimately conditional upon, is not, in general, an a priori matter. First, it is often an empirical question which desires we happen to have: Emmy might desire to develop lung cancer, she just happens not to. For her this might be an a priori matter, but for us it is not; and sometimes not even our own desires are a priori matters for us. Second, it is often an empirical question which means-end relationships obtain: that not smoking is a (non-deterministic causal) means to avoiding the development of lung cancer is an empirical, a posteriori discovery. In contrast to this, categorical imperatives do not depend on empirical matters that can only be discovered a posteriori. At least this is so if they are “grounded” in moral facts or other ontologically dubious categories that are not subject to empirical investigation in the way means-end relationships are.

Since Carnap (1963b: 983) holds that “questions of rationality are purely a priori,” he cannot understand norms as hypothetical imperatives. Assuming that

there is no third category of imperatives, it follows that Carnap must understand norms as categorical imperatives. Since the validity, or invalidity, of categorical imperatives cannot, in general, be discovered empirically, it is then only natural for him to postulate a powerful source of knowledge such as “rational insight” or “intuitive judgment” that allegedly enables us to know a priori which categorical imperatives hold.

Unfortunately philosophical problems are not solved by merely postulating solutions to them, no matter how many times we call them “rational.” That’s not engaging with a problem, that’s just giving up and not admitting it. Carnap, despite his empiricist confessions (Carnap & Hahn & Neurath 1929), seems to have indulged in just such a rationalist view of rationality. He seems to have thought of the conditions of adequacy as categorical imperatives: “*c* must define a system of bets that is [strictly] fair”. This, so it seems to be Carnap’s view, is a categorical imperative about how we should use the term ‘degree of confirmation,’ about how to explicate ‘degree of confirmation.’ On this view, we find out about the validity of this categorical imperative by intuitive judgment, which enables us to know a priori that the condition of adequacy holds.

Carnap’s curious view results because he seems to have thought of these *conditions of adequacy* as *categorical* imperatives of *explication* which justify the analytic truth of the axioms of inductive logic. In contrast to this, the view adopted in this paper holds that the *axioms* of inductive logic are *hypothetical* imperatives of *belief* that we justify by showing, i.e. providing an argument for the conclusion that, they are the means to attaining the cognitive end they are conditional upon.

Specifically, Carnap seems to have thought that conditions of adequacy are categorical imperatives that are formulated in the metalanguage and that tell us how we ought to use words in the object language. They thus take the form: it ought to be the case that *C*, $\mathbf{O}(C)$. The axioms of inductive logic are analytic truths in the object language of the form: *A*. The conditions of adequacy regulate the use of words in the object language. They are the reasons why we ought to accept the axioms. We thus get the following picture: $\mathbf{O}(C) \rightarrow A$, where \rightarrow is the reason (for acceptance) relationship between synthetic norms in the metalanguage and analytic truths in the object language.

On the present view the principle of induction and the axioms of inductive logic are not categorical imperatives to use words in certain ways, but hypothetical imperatives for organizing one’s beliefs *B* at a given moment in time. We justify these hypothetical imperatives to believe in a certain way by showing or arguing that they are the means to attaining the cognitive end *E* they are conditional upon. On the present view, the picture is this: $\mathbf{O}_E(B)$.

Which ultimate cognitive ends E we have is a factual question, or a matter of stipulation if we consider how we should believe if we had certain cognitive ends. The same is true for what we actually believe. Which means-end relationships obtain between one's believing, or intending to believe, B , and various ends E sometimes is an a priori matter, if theorem proving is – though not always, as we have seen above. For instance, to the extent that theorem proving is an a priori matter², we will see in section 5 that it can sometimes be discovered a priori that B -ing is a necessary, or a necessary and sufficient, condition for attaining E . (Whether this is also true for the mere intention to B depends on whether it is a priori that intentions are necessary for actions.) However, when the means-end relationship can be so established, it is by honest theorem proving, not by intuition mongering, rational or otherwise.³

4 The Continuum of Carnap's Rationalism

It is natural to wonder if Carnap was aware that his view commits him to a multitude of synthetic truths that we can know a priori, something whose existence he had denied earlier. This does not seem to be the case.

²I do not assume that theorem proving is an a priori matter, and nothing in this paper depends on the assumption that means-end relationships can be established a priori.

³There is another reason why Carnapian explication requires intuitive judgment to be a powerful source of a priori knowledge. According to Carnap (1962: 3) an explication is “the transformation of an inexact, prescientific concept, the *explicandum*, into a new exact concept, the *explicatum*.” An explicatum is to satisfy four requirements: similarity (in meaning), exactness, fruitfulness, and simplicity. “The explicatum is to be *similar to the explicandum* in such a way that, in most cases in which the explicandum has so far been used, the explicatum can be used; however, close similarity is not required, and considerable differences are permitted” (Carnap 1962: 7). In this respect explication differs from analysis, where the analysans is to be identical (in meaning) to the analysandum in such a way that, in all cases in which the analysandum has been used so far, the analysans can be used; strict identity (in meaning) is required, and no differences are permitted.

Explication requires intuitive judgment to be a powerful source of a priori knowledge, because the method of counterexamples does not work for explication. We can grant that Gettier (1963) shows that justified true belief is not identical in meaning to knowledge. What Gettier (1963) does not show is that justified true belief is not similar in meaning to knowledge, because a few counterexamples to the identity of meaning do not establish gross dissimilarity. Intuitive judgment already has to be a powerful source of knowledge if it can yield a priori knowledge that analysans and analysandum are not identical (in meaning). It has to be an even more powerful source of knowledge if it can yield a priori knowledge of gross dissimilarity (in meaning) between explicatum and explicandum.

In Carnap (1968) and in the unpublished manuscript “How can induction be justified?” in which he responds to Lenz (1956), and which seems to have served as a precursor of Carnap (1968), Carnap makes it very clear that he thinks he does not rely on any synthetic principles.⁴ Carnap distinguishes between two problems of justification. The second problem arises only if there is more than one confirmation function that satisfies the axioms of inductive logic, as is the case for the system presented in Carnap (1952). It is a problem of riches, namely the problem of justifying the choice of one of these confirmation functions over the others. The first problem is the problem of justifying the axioms of inductive logic themselves. It is this problem that corresponds to the problem of the justification of induction in Carnap’s work, and that is of relevance to the present paper. The story that Carnap (1968) tells is the same story as above, except that the reasons or conditions of adequacy are now called “requirements of rationality” (Carnap 1968: 260), and that the source of our a priori knowledge of these requirements of rationality is called “inductive intuition” (Carnap 1968: 265).⁵

Here is what Carnap (1968: 264) says about these requirements of rationality:

Now let us think about the question of the kind of factors we may legitimately appeal to in such arguments about rationality. As a matter of fact, in the history of the subject the following factors have often been used by one or the other philosopher:

- (a) deductive logic or deductive reasoning;
- (b) inductive reasoning;
- (c) past experience;
- (d) general synthetic principles (as for instance the principle of the uniformity of the world).

To save time, I will not speak about the last point, because I believe that today the majority of philosophers reject this, though some, for instance Russell, still appeal to such a principle.

Apparently Carnap is of the opinion that he does not rely on general synthetic principles (d). However, his very own definition of analyticity from the 1930s and

⁴I am grateful to André W. Carus for pointing me to, and providing me with, these three papers.

⁵In the unpublished manuscript Carnap also talks about “inductive inertia”, which plays a role for the second problem of riches, and which seems to be similar to what contemporary decision theorists call “risk” (Buchak 2014).

1940s, and his claim that “it is impossible to give a purely deductive justification of induction” from the 1960s imply that he does rely on such general synthetic principles. This will be clear once we recall that, for Carnap, inductive reasoning does not really differ from deductive reasoning, as both (a) and (b) consist of analytic truths.

Carnap (1968: 265) goes on to explain that both (a) and (b) are “legitimate” and “indispensable”, and “that the situation in deductive logic is exactly the same”: just as “deductive intuition” (Carnap 1968: 266) is the source of our a priori knowledge of the validity of deductive reasoning, so inductive intuition is the source of our a priori knowledge of the validity of inductive reasoning. Carnap (1968: 265) is aware that this “looks like defending a vicious circle,” but he

believe[s] all procedures of self-clarification, of making clear to ourselves what it is we have in mind, are in a way circular. We clarify *B* through *A*, and then we turn around and explain *A* with the help of *B*. I think we cannot do without it. If a person were unable to distinguish valid from invalid steps in inductive reasoning, even in the simplest cases, in other words, if he were inductively blind, then it would be hopeless to try to convince him of anything in inductive logic.

This quote makes clear that the problem Carnap is concerned with is how we can justify the results of explication⁶, of “making clear to ourselves what it is we have in mind.” Carnap is concerned with the justification of the axioms of (deductive and) inductive logic which govern the use of the term ‘degree of confirmation.’ In other words, Carnap is concerned with the justification of how we should *talk* about believing or reasoning. This is in sharp contrast with what I interpret Hume (1739; 1748) to be concerned with, and with what I am concerned with, namely the justification of (deductive and) inductive logic itself: the question why we should believe or reason in the way (deductive and) inductive logic prescribes.

Much like Carnap I will conclude in sections 6 and 7 that induction and deduction can be justified. However, in contrast to Carnap, I do not think that we can justify deduction and induction circularly, and by appeals to intuition in the way he has envisaged.

⁶Carnap (1962: 172f) distinguishes between explication and clarification. The latter is the process of clarifying the explicandum before it is explicated by an explicatum in explication. While the definition in an explication must not be circular, clarification indispensably is. I fail to see how clarification can be separated from explication, and Carnap (1968) does not seem to put much emphasis on this distinction anymore either.

In part this difference may be due to me thinking of deductive and inductive logic as normative theories regulating how we should believe or reason rather than thinking of the requirements of rationality as regulating how we should talk about believing or reasoning. In part this difference may also be due to me ascribing at best heuristic value to deductive and inductive intuition. Mainly, though, this difference is due to me not believing in categorical imperatives. In my opinion Carnap's view of rationality, or normativity, is wrong. Since he has the wrong theory of normativity, it is not surprising that Carnap also gets wrong the justification of normative theories, such as deductive and inductive logic.

5 Goodman's Fiction about the Facts of Forecast

Goodman (1946; 1947; 1954) is very critical of Carnap's syntactic program. The relevant section of the third chapter of *Fact, Fiction, Forecast* (1954) on "The New Riddle of Induction" is one of the best known pieces of 20th century philosophy of induction (see Stalker 1994).

Interestingly, though, the views expressed by Goodman in this chapter are very similar to Carnap's views on explication. Goodman follows Carnap in comparing inductive logic to deductive logic. He thinks that a particular deductive inference such as

Athens is the capital of Greece. Therefore, there is a capital of Greece.

is justified by conforming to the valid general rules of deductive logic, including

$$F(a) \therefore \exists x F(x)$$

The general rules of deductive logic in turn are justified by conforming to accepted deductive practice. Goodman (1983: 64) is, of course, aware that

[t]his looks flagrantly circular. ... But this circle is a virtuous one. The point is that 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.* The process of justification is the delicate one of making mutual adjustments between rules and accepted inferences; and in the agreement achieved lies the only justification needed for either.

At least to my ears this sounds very similar to Carnap (1968: 265)'s defense of the circular nature of "self-clarification." It also does not work. More specifically, it does not work for the deductive case, nor does it work for the inductive case.⁷

It is simply a fiction that it would be a fact that there is agreement between the particular deductive inferences we accept, and the general rules of deduction that are valid. Many do not accept particular deductive inferences such as

Athens is the capital of Greece. Therefore, Athens is the capital of Greece or the moon is made of cheese.

This particular deductive inference is not accepted by many of us, I claim, though only few have attempted to amend the following general rule of deduction that is valid (in classical logic) and that the particular deductive inference conforms to:

$A \therefore A \vee C$

Many more examples can be given, even if we bracket the material conditional whose truth conditions presumably have caused incredulous stares on the faces of most of us when we first encountered them. Indeed, far from there being any agreement, the discrepancy between the particular deductive inferences we accept, and the general rules of deduction that are valid, is *so* large that some philosophers have even felt the need for supplementary theories to bridge the gap. Grice (1975) proposes a supplementary theory according to which the maxims of quantity and of relevance and of manner that require us to be informative and relevant and perspicuous, respectively, explain why many do not accept the particular deductive inference above, even though it conforms to a general rule of deductive logic that is valid. The latter rules matter mainly for the maxim of quality according to which we should tell the truth. Jackson (1979) proposes a supplementary theory that attempts to explain the gap between particular accepted inferences involving indicative conditionals and general valid rules of deductive logic involving the material conditional. Lewis (1986: fn 6) adopts this theory. Apparently all these philosophers agree that the discrepancy between accepted deductive inferences and valid deductive rules is so large that a supplementary theory is needed.

The problems Goodman's view has with the deductive case carry over to the inductive case. Indeed, the very difference between deductive and inductive logic seems to vanish on Goodman's view. We have seen that the difference between deductive and inductive reasoning vanishes for Carnap: both his (a) and (b) consist of analytic truths.

⁷The same is true of Strawson (1952: ch. 9; 1958)'s attempt to define away the problem of induction.

The situation is somewhat similar on Goodman's view. There are particular inferences some of us accept, and other particular inferences those same of us do not accept. On Goodman's view it seems arbitrary to further divide the acceptable particular inferences into those that are acceptable *qua* deductive inferences, and those that are acceptable *qua* inductive inferences. The inference from Athens being the capital of Greece to there being a capital of Greece is just as acceptable as a deductive inference as it is acceptable as an inductive inference. The inference from Athens being the capital of Greece to Athens being the capital of Greece or the moon being made of cheese is just as unacceptable as a deductive inference as it is unacceptable as an inductive inference.⁸

To be sure, there is a difference between deductive and inductive logic. However, this difference is not grounded in the source of our logical knowledge, the difference between deductive and inductive intuition, as Carnap has it. Nor is this difference grounded in our inferential practice, the difference between inferences that are acceptable *qua* deductive inference and inferences that are acceptable *qua* inductive inferences, as Goodman will have it. The difference between deductive and inductive logic is a difference in the cognitive ends these normative systems are a means to. In this sense this difference is grounded in the normativity of logic.

How do things work in the deductive case? How do we justify the general rules of deductive logic? We justify them in the exact same way as we justify all norms, namely by showing or arguing that they are the means to some end (whether this end is desired by anybody matters only for the relevance of the norm in question, not its validity). The general rules of deductive logic are justified relative to the end of reasoning in a way that is truth-preserving with logical necessity. For all and only those particular inferences that conform to the general rules of deductive logic are truth-preserving with logical necessity. The epistemological significance of Gödel (1930)'s correctness and completeness theorems for classical logic lies in establishing this means-end relationship. The epistemological significance of Gödel (1931)'s (first) incompleteness theorem lies in establishing that, with logical necessity, we miss out on something regarding the end of holding only true beliefs, and as many thereof as possible, once we widen our cognitive interests and become curious about arithmetical questions (and believe an answer to an arithmetical question just in case we can derive it in Peano arithmetic). For either not every statement of the language of Peano arithmetic that is derivable in Peano

⁸Goodman's positive theory of induction in terms of "entrenchment" is developed in Goodman (1983: sct. IV). If (inductive) logic is a normative theory it says that we ought to project or infer inductively what we do in fact project or infer inductively, and thus infers an *Ought* from an *Is*.

arithmetic is true (in the standard model) – in this case we miss out by not holding only true beliefs; or not every statement of the language of Peano arithmetic that is true (in the standard model) is derivable in Peano arithmetic – in this case we miss out by not holding as many true beliefs as possible.

In section 7 I will turn to the entirely valid objection that this justification of the rules of deductive logic relative to the end of reasoning in a way that is truth-preserving with logical necessity is viciously circular, because it uses the rules of deductive logic. Before doing so, let me briefly give two more examples of how we justify a system of logic understood as a system of hypothetical imperatives governing how one should believe or reason relative to a particular cognitive end.

The first example is the non-monotonic logic developed by Kraus & Lehmann & Magidor (1990). Its rules are justified – not relative to the end of reasoning in a way that is truth-preserving with logical necessity, but – relative to the end of reasoning in a normally truth-preserving way: all and only those particular inferences that obey these rules are such that their conclusion is true in all of the most normal possible worlds in which all premises are true. Three things are interesting about this example. First, I think this system has a better chance to come out as the system of deductive logic on Goodman's approach, even though it agrees with classical deductive logic on the examples above. Second, this system is just as much a deductive logic as it is an inductive logic, thus stressing that this dichotomy is somewhat arbitrary. Third, this example illustrates (as does the next) that sometimes it can be discovered a priori (see, however, footnote 2) that obeying a system of norms is a provably necessary and sufficient means to attaining a particular end. However, it is by honest theorem proving rather than by appeals to intuition that we can obtain such a priori knowledge.

The first example illustrates that different cognitive ends give rise to different logical systems. The next example illustrates that different propositional attitudes give rise to different logical systems. It is the probability calculus whose rules are justified relative to the end of avoiding accuracy domination (Joyce 1998; 2009): all and only those degree of belief functions that obey these obligations are such that it is not the case that they are accuracy dominated in the sense that there exists an alternative degree of belief function that is at least as accurate in all possible worlds, and more accurate in at least one. This example is an instance of epistemic consequentialism (Percival 2002, Stalnaker 2002, Brössel & Eder & Huber 2013) according to which epistemic rationality is a species of instrumental rationality. If we supplement this view with explicitly denying the existence of objective values so that there is nothing objectively wrong or right about what ultimate cognitive ends one happens to have we arrive at the view adopted here.

The upshot of this is three-fold. First, we need to specify a cognitive end before we can even meaningfully ask if we can justify the principle of induction relative to this end. We have already seen that the problem of induction vanishes for a cognitively contented being. Now we see that different cognitive ends give rise to different principles of induction, or inductive logics, and different problems of induction.

Second, we must pay attention to the fact that different propositional attitudes require different principles of induction. Not all cognitive ends are meaningful for all propositional attitudes. For instance, someone who has only degrees of belief might not be able to meaningfully desire the cognitive end of holding only true beliefs, and as many thereof as possible.

Third, we must not forget that different agents have different cognitive abilities and limitations. If our doxastic agent is omniscient, she does not need a principle of induction. If someone sometimes forgets, is not always able to voluntarily decide what to believe (to a specific degree), or cannot compute everything computable, our principle of induction needs to take these limitations into account.

6 Reichenbach Rules

All this seems to have been clear to Reichenbach (1938; 1940) who was interested in figuring out the limit of the relative frequency of types of events in sequences of events. To this end he formulated the straight rule, which is his version of the principle of induction, his inductive logic. Here is my preferred variant of the straight rule: from the premise that m out of the n events about which one has the information whether they are of type A are of type A , one may and ought to infer the conclusion that the limit of the relative frequency of events of type A in any sequence of events continuing these n events is m/n .

Whatever one thinks of the end of figuring out limits of relative frequencies of types of events in sequences of events, it is insightful to consider how Reichenbach justifies the straight rule relative to this end. He does so by providing a deductively valid argument which does not presuppose its conclusion, and whose premises contain enough mathematics to make sense of the numbers that are mentioned in the straight rule. The conclusion of this deductively valid argument does not say that the straight rule usually leads from true premises to true conclusions. Instead it says that obeying the straight rule is a means to attaining the cognitive end of converging to the true limit of the relative frequency of any type of event in any sequence of events for which there is such a limit.

What Reichenbach shows is that, for any sequence of events s and any type of event A : the limit of the relative frequency of A in s exists if and only if the straight rule converges to it. Logicians may want to think of this result as establishing the correctness and completeness of Reichenbach's inductive logic with respect to the cognitive end of converging to the true limit of the relative frequency of any type of event in any sequence of events for which there is such a limit: provided the information in the premises is true, the straight rule converges only to true limits (correctness), and it converges to all true limits (completeness).

One of the objections to Reichenbach's argument is that the straight rule is not the only rule that converges to all and only true limits. The rule that conjectures $(m + 5) / n$ when the straight rule conjectures m/n does so, too. In response to this objection one may point out that the situation is exactly analogous in deductive logic. There are many different axiomatizations of classical logic. However, all these axiomatizations are *functionally equivalent* in the sense that all of them characterize the exact same set of particular inferences. This means that, as far as the cognitive end of reasoning in a way that is truth-preserving with logical necessity is concerned, all these different axiomatizations are on a par. Note that these axiomatizations are not on a par with respect to other ends – say, the end of deriving conclusions in as few steps as possible.

Consider a different example. Drinking a glass of water quenches your thirst if and only if drinking a glass of water or some other drink does. As far as the end of quenching your thirst is concerned drinking a glass of water is on a par with drinking another drink. The two are functionally equivalent with respect to this end. Of course, there are other ends, such as quenching your thirst with a sweet drink or a drink with few calories, with respect to which drinking a glass of water is not on a par with drinking another drink. But that is besides the point, because that is not claimed. Similarly, the straight rule and the above alternative are on a par as far as the cognitive end of converging to the true limit of the relative frequency of any type of event in any sequence of events for which there is such a limit is concerned. The two are functionally equivalent with respect to this cognitive end. It is not claimed that these two rules are functionally equivalent with respect to any other end – say, making true or accurate conjectures about the type of the next event.

Given that you have the end of quenching your thirst with a drink with few calories, you ought to drink a glass of water. Yet you cannot *just* drink a glass of water. You can only drink a particular glass of water containing a certain number of H_2O molecules at a particular time in a particular manner. However, drinking another glass of water containing the same number of H_2O molecules, but differ-

ent ones, at a slightly later time in a slightly different manner would do equally well: it would be functionally equivalent with respect to the end of quenching your thirst with a drink with few calories. Therefore, given that you have this end, what you really ought to do is something that is functionally equivalent to drinking a particular glass of water containing a certain number of H₂O molecules at a particular time in a particular manner, where functional equivalence is understood with respect to the end of quenching your thirst with a drink with few calories. In the same way we cannot *just* follow the straight(-forward) rule or the above alternative or some other alternative that converges to the same limits as the straight rule. We can only follow *one* of these alternative rules. Therefore, all that the non-uniqueness objection to Reichenbach's argument shows is that we have to formulate its conclusion more carefully as follows: given that we have the cognitive end of converging to the true limit of the relative frequency of any type of event in any sequence of events for which there is such a limit, we ought to *obey a rule that is functionally equivalent to the straight rule*, where functional equivalence is understood with respect to this cognitive end.

Another objection is the slogan, associated with Keynes, that "in the long run we are all dead." This suggests that nobody has, or should have, the end of converging to the true limit of the relative frequency of any type of event in any sequence of events for which there is such a limit. Even if this was the case, the objection would miss the point. Reichenbach's argument establishes a means-end relationship without suggesting that one ought to, or does, have the relevant end.

Furthermore, the situation in deductive logic is again analogous. Just as in the long run we are all dead, nobody lives in any logically possible world other than the actual one. Yet this does not prevent people from desiring to have beliefs that are consistent, i.e. true in *some* logically possible world that need not be the actual one. Now, consistency alone may not make beliefs desirable, but think about how undesirable beliefs are that are not even consistent. Similarly, convergence to the true limit alone may not make a sequence of conjectures desirable, but think about how undesirable a sequence of conjectures is that does not even converge to true limit. It is no coincidence that estimators in statistics are called "strongly consistent" if their conjectures converge (with probability one, or almost surely).

A third objection is the reference class problem. In fact, there are two of these problems: one for the type of event *A* (each event instantiates many types) and one for the sequence of events *s* (each event belongs to many sequences). Picking the narrowest reference class for *A* (the type whose only instantiation is the next event) is picking a type that contains too few events, and picking the broadest reference class for *s* (the sequence of all events) is picking a sequence

that contains too many events. Picking any other reference class is arbitrary.

The reference class problems are indeed problems if one wants to use observed relative frequencies to determine probabilities, say, the probability that the next event is of a certain type. However, they are *not* problems for the straight rule, because the latter does not attempt to do this. The straight rule does not even mention probability or the next event. Instead it conjectures the limit of the relative frequency of any type of event A in any sequence of events that continues the n events about which one has the information whether they are of type A .

Now, to be sure, Reichenbach was *also* interested in using observed relative frequencies to determine the probability that the next event is of a certain type. To this end he introduced the notion of a *posit*. The justification of the straight rule with respect to the end of converging to all and only true limits does not carry over to this different end. However, let us not throw out the baby with the bathwater. Reichenbach taught us that we can justify induction deductively by proving it to be the means to attaining an end that differs from usually reasoning true premises to true conclusions.

Specifically, we can deductively justify the principle of universal induction relative to the cognitive end of necessarily “stabilizing” on all and only adequate universal hypotheses, i.e. hypotheses of the form ‘All F s are G .’ For there is a deductively valid argument which does not presuppose its conclusion, whose premises are restricted to some definitions, and whose conclusion says the following: in each possible world w , the principle of universal induction stabilizes on all and only those universal hypotheses that are adequate in w . Here I identify a possible world with an ordered set of informational inputs. The talk of adequacy, as opposed to truth, is intended to stress this assumption. I also assume that the information we have at any point in any possible world w is true in w as well as “cumulative” in a sense that will be made precise momentarily.

The premises of this deductively valid argument can be restricted to some definitions because, unlike Reichenbach’s straight rule, the principle of universal induction does not mention numbers or anything else that requires more than first-order logic with identity. As to the definitions, they are needed to formulate the end which the principle of universal induction is a means to attaining, and which is mentioned in the conclusion of this deductively valid argument. The latter says that for all possible worlds w and all universal hypotheses h , h is adequate in w if and only if there is a point such that at this and any later point, h may and ought to be inferred from the information one has at these points in w by the principle of universal induction.

Let’s formulate this conclusion more precisely with some definitions. First

I assume there to be a set of points P whose elements are ordered by a binary relation \leq that is anti-symmetric, transitive, and total: for all s, t, u in P , if $s \leq t$ and $t \leq s$, then $s = t$; if $s \leq t$ and $t \leq u$, then $s \leq u$; and, $s \leq t$ or $t \leq s$. Possible worlds are characterized as functions w from P into the set of literals L , i.e. atomic sentences and their negations, over some first-order vocabulary such that for all atomic formulas $F(a)$: $F(a)$ is in the range of w if and only if $\neg F(a)$ is not. The range of w is the set of literals $\pm F(a)$ in L for which there is a point t in P such that $w(t) = \pm F(a)$. The universal hypothesis ‘All F s are G ’ is adequate in the possible world $w : P \rightarrow L$ if and only if there is no individual constant a such that $F(a) \wedge \neg G(a)$ follows logically from the range of w , $I(w) = \{w(t) : t \in P\}$. Otherwise it is inadequate in w .

The information one has at some point t in P in a possible world w is the set of all literals $w(s)$ such that $s \leq t$, $I(w, t) = \{w(s) : s \in P, s \leq t\}$. This definition implies that the information one has at any point t in any world w is true in w : $I(w, t) \subseteq I(w)$. It also follows that every bit of information that is true in w will be received eventually: if some item of information belongs to $I(w)$, then it belongs to $I(w, t)$, for some t in P . Finally, once received, an item of information will never be forgotten. This is so because the information one has at any point in any possible world is *cumulative*: for all points s and t , and all possible worlds w , if $s \leq t$, then $I(w, s) \subseteq I(w, t)$.

Let Def be a set of first-order formulas describing all these definitions. This means that we cannot make the assumptions that P is countable and that \leq is a well-ordering, because countability and well-ordering cannot be formulated in first-order logic. As mentioned, Def is needed to formulate the *end* which the principle of universal induction is a means to attaining. Def is not needed to formulate the principle of universal induction itself. The following proof can be carried out in first-order logic with identity.

Let w be an arbitrary possible world and h an arbitrary universal hypothesis, i.e. statement of the form ‘All F s are G ,’ where ‘ F ’ and ‘ G ’ are arbitrary predicates (and not just atomic ones). h is adequate in w , or h is inadequate in w . If h is adequate in w , then there is no individual constant a such that $F(a) \wedge \neg G(a)$ follows logically from $I(w)$. Since $I(w, t) \subseteq I(w)$ for all points t , there is no individual constant a such that $F(a) \wedge \neg G(a)$ follows logically from $I(w, t)$, for any point t . Hence at any point t in w , all objects, if any, about which one has enough information – i.e. whether they are F and whether they are G – are G if they are F . According to the principle of universal induction, from all these premises $I(w, t)$ one may and ought to infer the conclusion that all F s are G .

If h is inadequate in w , then there is an individual constant a such that $F(a) \wedge$

$\neg G(a)$ follows logically from $I(w)$. By compactness there is a finite subset $I(w)^-$ of $I(w)$ that logically implies $F(a) \wedge \neg G(a)$. Since $I(w)^-$ is finite, there exists a point s such that $t \leq s$ for all points t with $w(t) \in I(w)^-$. Since information is cumulative it holds for all points u such that $s \leq u$: *not* all objects about which one has enough information – i.e. whether they are F and whether they are G – at point u in w are G if they are F . Therefore it holds for all points u such that $s \leq u$: it is not the case that one may and ought to infer the conclusion that all F s are G from the premise $I(w, u)$ by the principle of universal induction.

The universal hypotheses that the principle of universal induction conjectures will, in general, be “over the top.” For instance, as long as there is no object about which one has the information whether it is F it conjectures that all F s are $\neg F$. This boldness is necessary for conjecturing all adequate universal hypotheses “as quickly as possible.” Before showing this let us put this first result in symbols, and in various flavors that are not logically equivalent, but true nonetheless:

$$\begin{aligned} \text{Def } \vdash_{DL} \forall w \forall h (w \models h \rightarrow \exists s \forall t \geq s : I(w, t) \vdash_{PoI} h) \\ \text{Def } \vdash_{DL} \forall w \forall h (w \not\models h \rightarrow \exists s \forall t \geq s : I(w, t) \not\vdash_{PoI} h) \\ \text{Def } \vdash_{DL} \forall w \forall h (w \models h \leftrightarrow \exists s \forall t \geq s : I(w, t) \vdash_{PoI} h) \\ \text{Def } \vdash_{DL} \forall w \forall h (w \not\models h \leftrightarrow \exists s \forall t \geq s : I(w, t) \not\vdash_{PoI} h) \end{aligned}$$

\vdash_{DL} is derivability in deductive logic DL , and \vdash_{PoI} is derivability from the principle of universal induction PoI . \models is the relation of adequacy, and the variables ‘ h ’ and ‘ w ’ range over universal hypotheses and possible worlds, respectively. The variables ‘ s ’ and ‘ t ’ range over points in P , and $t \geq s$ holds if and only if $s \leq t$. The following formula says that there is a derivation in first-order logic with identity from our definitions to the claim that, with necessity, the principle of universal induction stabilizes on all and only adequate hypotheses: $\text{Def } \vdash_{DL} \text{NecSt}(PoI)$.

Logicians may want to think of this result in terms of correctness and completeness: the principle of universal induction stabilizes *only* on adequate universal hypotheses (correctness), and it stabilizes on *all* adequate universal hypotheses (completeness). This means that reasoning in accordance with the principle of universal induction is a means to attaining the cognitive end of necessarily stabilizing on all and only adequate universal hypotheses.

As in the case of the straight rule, this result does not uniquely characterize the principle of universal induction. Uniqueness is achieved by focusing on a more ambitious cognitive end. Among all principles that have the same abilities and limitations as PoI (effectively, that have no computational limitations and, at point t in possible world w , have access to, and only to, $I(w, t)$), the latter is such

that in all possible worlds w : PoI is at least as fast in w as any other principle when it comes to inferring universal hypotheses that are adequate in w . This is so, because PoI conjectures that all F s are G , for any predicates ‘ F ’ and ‘ G ,’ at any point t at which the information one has at t in any world w does not logically imply $F(a) \wedge \neg G(a)$ for some individual constant a . In particular, this is so when one does not even have any information about any F s or G s – or, indeed, any information at all. Therefore no principle is faster in any possible world when it comes to inferring adequate universal hypotheses (as well as inadequate ones).

This more ambitious cognitive end still leaves us with several principles that are all functionally equivalent in the sense of necessarily stabilizing on all and only adequate universal hypotheses, and inferring adequate universal hypotheses as quickly as possible. However, among the latter principles PoI is uniquely characterized by being the one and only principle that *additionally* avoids inadequate universal hypotheses as quickly as possible. For suppose there is a principle PoJ and a possible world w and a universal hypothesis h of the form ‘All F s are G ’ that is inadequate in w , and PoJ stops inferring h in w at some point t , *before* PoI does. Consider a possible world w^* that is identical to w in all points s with $s \leq t$, but that differs from w insofar as h is adequate in w^* . Such a possible world w^* exists, because PoI still conjectures that all F s are G at t in w , and so no $I(w, s)$, for any $s \leq t$, logically implies $F(a) \wedge \neg G(a)$ for any individual constant a . Since PoJ has the same abilities and limitations as PoI , PoJ still – but wrongly! – stops inferring h in this alternative possible world w^* at point t . If PoJ never resumes inferring h in w^* , or does so, but without stabilizing on h , then PoJ does not stabilize on all universal hypotheses that are adequate in w^* . If PoJ resumes inferring h in w^* and stabilizes on h , then it does not do so as fast as PoI . In either case PoJ is not one of those principles that necessarily stabilize on all and only adequate universal hypotheses, and infer adequate universal hypotheses as quickly as possible.

Thus the principle of universal induction is the one and only means to attaining the cognitive end of necessarily stabilizing on all and only adequate universal hypotheses, of inferring adequate universal hypotheses as fast as possible, and of *subsequently* avoiding inadequate universal hypotheses as fast as possible. (I write ‘subsequently’ to stress that the former end has priority over the latter.)

The work by Reichenbach (1938; 1940) and others (Putnam 1963, Gold 1967) has inspired an entire discipline of deductive justifications of induction: *formal learning theory*. It is developed as brilliantly as it is presented beautifully in Kelly (1996) and Martin & Osherson (1998). For the purposes of this paper the variant of Reichenbach’s deductive justification of the principle of universal induction just given is all we need. Let us now turn to the inductive justification of deduction.

7 Haack's Argument against Deduction

Carnap (1968) and Goodman (1954) compare induction to deduction. Haack (1976) compares deduction to induction. Her preferred variant of Hume (1739; 1748)'s argument against induction takes the form of a dilemma.

Hume's dilemma for induction A deductive justification of induction would be too strong, because it would show that induction always leads from true premises to true conclusions.

An inductive justification of induction would be circular.

Haack's dilemma for deduction An inductive justification of deduction would be too weak, because it would merely show that deduction usually leads from true premises to true conclusions.

A deductive justification of deduction would be circular.

Haack (1976) is right in claiming that an inductive justification of induction would be circular, and that a deductive justification of deduction would be circular. These are two immediate circles that I will return to below. However, even if we grant Premise 1 of Hume's Argument, Haack (1976) is not right in claiming that a deductive justification of induction would be too strong. It is true, as we have seen, that there is no deductively valid argument which does not presuppose its conclusion, whose premises are restricted to information we have, and whose conclusion says that the principle of induction usually leads from true premises to true conclusions. And perhaps this is what Haack (1976) means. It is not true, though, that the latter conclusion would be too strong to justify the principle of induction.

It is also not true that an inductive justification of deduction would be too weak, or so I will argue now. We have the following piece of information:

All particular inferences which are such that we have the information (i) whether their ingredient sentences (i.e. premises plus conclusion) are true or false, and (ii) whether they conform to the general rules of classical deductive logic, have at least one false premise or a true conclusion if they conform to the general rules of deductive logic.

From this premise we get the following conclusion by an application of the principle of universal induction:

All particular inferences that conform to the general rules of classical deductive logic have at least one false premise or a true conclusion.

This conclusion says that classical deductive logic *always* leads from true premises to true conclusions. It does not merely say that classical deductive logic *usually* leads from true premises to true conclusions. Since the argument also does not presuppose its conclusion, and its premise is restricted to information we have, we have a perfectly fine inductive justification of deduction. Pace Haack (1976) it is not the case that an inductive justification of deduction would be too weak.

More precisely, we have an inductive justification of classical deductive logic relative to the cognitive end of reasoning in an *actually* truth-preserving way. What we do not have is an inductive justification of classical deductive logic relative to the different cognitive end of reasoning in a way that is truth-preserving *with logical necessity*: the conclusion of our inductive justification of classical deductive logic does not say that the particular inferences conforming to the general rules of classical deductive logic lead from true premises to true conclusions *with logical necessity*. And perhaps this is what Haack (1976) means when she says that an inductive justification of deduction would be too *weak*. The conclusion merely says that particular inferences which conform to the general rules of classical deductive logic *actually* and *always* lead from true premises to true conclusions. However, the conclusion does not merely say that particular inferences which conform to the general rules of classical deductive logic *normally* or *usually* lead from true premises to true conclusions.

Even if this were the complaint, it is not the principle of induction that is to be blamed. If anyone, it is *we* who are to be blamed, because *our* cognitive limitations prevent us from establishing the premise needed to inductively infer this stronger conclusion without presupposing it. Suppose we had the following information:

All particular inferences which are such that we have the information (i) whether their ingredient sentences (i.e. premises plus conclusion) are true or false, and (ii) whether they conform to the general rules of classical deductive logic, are such that, if they conform to the general rules of deductive logic, then, with logical necessity, at least one of their premises is false or their conclusion is true.

From this premise we would indeed get the following stronger conclusion by an application of the principle of universal induction:

All particular inferences that conform to the general rules of classical deductive logic are such that, with logical necessity, at least one of their premises is false or their conclusion is true.

The problem is that *we* cannot show that the particular inferences which are such that we have the information whether their ingredient sentences are true or false, and that conform to the general rules of classical deductive logic, are truth-preserving with logical necessity without presupposing that particular inferences that conform to the general rules of classical deductive logic are truth-preserving with logical necessity. The problem is *not* the principle of universal induction, which would happily yield this conclusion for us if our cognitive limitations did not prevent us from having as information the premise needed to inductively infer this stronger conclusion.

Besides being able to justify deduction inductively relative to the cognitive end of reasoning in an actually truth-preserving way, we saw in section 6 that we can also justify various principles of induction deductively relative to various cognitive ends (and, of course, inductively, as many philosophers have noted). Moreover, we can justify deduction deductively relative to the end of reasoning in a way that is truth-preserving with logical necessity, as we have seen in section 5. To be sure, this deductive justification of deduction is circular. However, it is neither trivial nor need this circularity concern us.

It is not trivial, because the rules of classical deductive logic do not justify themselves deductively relative to all ends (for instance, they do not justify themselves relative to the end of reasoning in a normally truth-preserving way). Nor do they justify the rules of other systems of logic (for instance, the non-monotonic logic by Kraus & Lehmann & Magidor 1990) with respect to the end of reasoning in a way that is truth-preserving with logical necessity. Nor do they justify themselves relative to the end of reasoning in a way that is truth-preserving with logical necessity once we enrich them to include further axioms.

The circularity need not concern us, because classical deductive logic can be justified inductively relative to the end of reasoning in an actually truth-preserving way. While other systems of logic can also justify themselves in a circular way relative to some end, none of these other systems of logic will be justified inductively relative to the end of reasoning in an actually truth-preserving way. This is so because of our previously established result that *all* general rules of classical deductive logic are actually truth-preserving *and*, as we will see now, because *only* the general rules of classical deductive logic are actually truth-preserving.

The latter is the case for the following reason: any system of logic that contains at least one general rule that is not already part of classical deductive logic will be such that there is at least one particular inference that conforms to this general rule and has true premises, but a false conclusion. (Logicians call rules that go beyond classical deductive logic *inadmissible* in it. They call the rules *schemata*. And

they point out that the rules or schemata must not mention any logical constants other than those of classical first-order logic with identity.) Call a general rule of logic *tested* by us if and only if there is at least one particular inference of which we have the information whether its ingredient sentences are true or false, and that it conforms to this rule. I assume we have the following information:

For any general rule of logic that we have tested, and of which we have the information whether it goes beyond classical deductive logic, if it goes beyond classical logic, then there is at least one particular inference that conforms to this rule, whose premises are true, and whose conclusion is false.

What I am assuming here is that we can provide a “counterexample” to any rule that we have tested and of which we have the information that it is inadmissible in classical logic. That is, we can provide a particular inference of which we have the information that it conforms to the rule, that its premises are true, and that its conclusion is false. An application of the principle of universal induction to the above piece of information gives us the following conclusion:

For all general rules of logic that go beyond classical deductive logic there is at least one particular inference that conforms to this rule, whose premises are true, and whose conclusion is false.

Logicians may want to think of this conclusion as a completeness result (*only* the general rules of classical deductive logic are actually truth-preserving), and of the previous one as a correctness result (*all* general rules of classical deductive logic are actually truth-preserving).

Consequently we have an inductive justification of the general rules of classical deductive logic relative to the cognitive end of reasoning in an actually truth-preserving way: there is an inductively strong argument which does not presuppose its conclusion, whose premises are restricted to information we have, and whose conclusion says that all and only the general rules of classical deductive logic are actually truth-preserving. Here a general rule of logic is actually truth-preserving just in case all the particular inferences conforming to it have at least one false premise or a true conclusion. In symbols (and combining $\forall x (Ax \rightarrow Bx)$ and $\forall x (Bx \rightarrow Ax)$ to $\forall x (Ax \leftrightarrow Bx)$ for the sake of perspicuity):

$$I(@, now) \vdash_{PoI} \forall R (TP_{@}(R) \leftrightarrow Adm_{DL}(R))$$

$$I(@, now) \vdash_{PoI} \forall R (\forall i_R : @ \models i_R \leftrightarrow \forall i_R : \vdash_{DL} i_R)$$

$TP_{@}$ is actual truth-preservation, and Adm_{DL} is admissibility in deductive logic DL . \models now is the relation of truth, and $I(@, now)$ is the information we have now in the actual world. The variables ‘ R ’ and ‘ i_R ’ range over rules and instances, respectively. An instance is identified with a disjunction whose disjuncts are the negations of all premises, and the conclusion. The following formula says that there is an application of the principle of universal induction whose premise is the information we have now in the actual world and whose conclusion says that first-order logic with identity is actually truth-preserving: $I(@, now) \vdash_{PoI} ActTP(DL)$.

The inductive justification of classical deductive logic obtains relative to the end of reasoning in an actually truth-preserving way. The deductive justification of the principle of universal induction obtains relative to the different end of necessarily stabilizing on all and only adequate universal hypotheses (and, if one chooses to add so, of inferring adequate universal hypotheses as quickly as possible, and of subsequently avoiding inadequate universal hypotheses as quickly as possible). One might think I have replaced two immediate circles, the inductive justification of induction and the deductive justification of deduction, by one not so immediate, but still circular circle. This would be wrong, though.

It is not the case that we justify, in some logic L , deduction inductively, and induction deductively, let alone that we justify the package consisting of both induction and deduction in a deductive-or-inductive – or any other – way. Nor is it the case that we justify, in some logic L , the claim that deductive logic is (actually or necessarily) truth-preserving by assuming inductive logic to be (actually or necessarily) stabilizing on all and only adequate universal hypotheses, and the claim that inductive logic (actually or necessarily) stabilizes on all and only adequate universal hypotheses by assuming deductive logic to be (actually or necessarily) truth-preserving. That is, neither of the following is the case:

$$PoI \vdash_L DL \vdash_L PoI \quad \text{or} \quad DL \vdash_L PoI \vdash_L DL$$

$$DL \& PoI \vdash_L DL \& PoI \quad \text{or} \quad X \vdash_{DL \& PoI} DL \& PoI \quad \text{or} \quad X \vdash_L PoI \& DL$$

$$ActSt(PoI) \vdash_L ActTP(DL) \vdash_L ActSt(PoI)$$

$$ActSt(PoI) \vdash_L NecTP(DL) \vdash_L ActSt(PoI)$$

$$NecSt(PoI) \vdash_L ActTP(DL) \vdash_L NecSt(PoI)$$

$$NecSt(PoI) \vdash_L NecTP(DL) \vdash_L NecSt(PoI)$$

Rather, we use induction to work our way towards an inductive justification of deduction relative to one cognitive end. And we use deduction to work our way

towards a deductive justification of induction relative to a different cognitive end. As long as these two cognitive ends are different, we are not caught in a circle.

$$I(@, now) \vdash_{PoI} ActTP(DL) \quad \text{and} \quad Def \vdash_{DL} NecSt(PoI)$$

Maybe it helps to illustrate the situation with an analogy. Let me begin with a *disanalogy*. Ingrid and Desmond are two friends of mine. Ingrid tells me that all and only those claims are true that Desmond claims to be true. Desmond tells me that all and only those claims are true that Ingrid claims to be true. Suppose what Desmond tells me is true: all and only those claims are true that Ingrid claims to be true. Then Ingrid's claim is true that all and only those claims are true that Desmond claims to be true. And so Desmond's claim is true that all and only those claims are true that Ingrid claims to be true. Now suppose what Desmond tells me is false. Then Ingrid's claim is false that all and only those claims are true that Desmond claims to be true. And so Desmond's claim is false that all and only those claims are true that Ingrid claims to be true. If this were our situation, we would indeed be caught in a circle.

However, this is not our situation. Our situation is more analogous to the following one. Ingrid and Desmond are two friends of mine. Their expertise lies in different, though overlapping areas, and while they speak the same language, the way they speak is very different. Ingrid refines her claims over time, and while her claims are restricted to the location of their utterance, they are most interesting. Desmond, on the other hand, only makes seemingly platitudinous claims and he has already said everything he had to say. While Desmond's claims were not very bold, they were claimed to be true everywhere. Ingrid's claims vary with time and location, Desmond's do not.

Ingrid has just reassured me (as she has always as far as I remember) that the following is true here: all and only those general rules are truth-preserving here that Desmond claimed to be truth-preserving everywhere.

Desmond has claimed that the following is true everywhere: wherever Ingrid has been, or will be, speaking, all and only those universal hypotheses are true there that Ingrid there eventually continues to claim to be true there forever.

In the first case Ingrid and Desmond praise each other of the same qualities in the same way. In this sense the first case is circular. In the second case Ingrid and Desmond also praise each other, but they praise each other of different qualities, and they praise each other in very different ways. In the first case Ingrid's claim is true just in case Desmond's claim is true, which just means that Ingrid's claim is false just in case Desmond's claim is false. This is circularity, but it is absent in the second case, which illustrates our situation.

Appendix

The following sentence schema S gives rise to a possible objection to the inductive completeness argument:

$$\exists x \neg \exists y P(y, x) \wedge \forall x \exists y P(x, y) \wedge \forall x \forall y \forall z (P(x, z) \wedge P(y, z) \rightarrow x = y)$$

No instance of S is true in the model-theoretic sense in any finite model. However, the instance of S that replaces ‘ P ’ with ‘the natural number ... is a precursor of the natural number ...’ is true. That is, S is true in the model whose domain is the set of natural numbers and whose interpretation function interprets ‘ P ’ as the precursor relation on the natural numbers. Therefore every instance of the negation of S , $\neg S$, is true in every finite model, but some instance of $\neg S$ is false in some infinite model. This means that the general rule of logic $\therefore \neg S$ is not valid.

Here is the objection. Presumably the proponent of the inductive completeness argument is an empiricist who holds that every particular inference of which *she*, i.e. the empiricist, has the information that it conforms to this general rule, and whether its ingredient sentences are true or false, has (at least one false premise or) a true conclusion. After all, how could an empiricist claim something that is only true in infinite models. Yet if this is so, then this general rule comes out as valid on the inductive completeness argument, even though it is not!

My response to this objection is that it applies a double standard. Either the objector has the information that the instance of S described above (or some other instance) is true, or else she does not. If, as her objection suggests, she does, then *she*, i.e. the objector, does *not* have the information that every particular inference of which *she*, i.e. the objector, has the information that it conforms to this general rule, and whether its ingredient sentences are true or false, has (at least one false premise or) a true conclusion. In this case the above rule does *not* come out as valid on the inductive completeness argument. If, on the other hand, the objector does not have the information that the instance of S described above (or some other instance) is true, then she does not have the information that the rule is not valid. In this case she cannot raise her objection.

Thus this objection to the inductive completeness argument rests on applying a double standard. The objector assumes to have information about arithmetical claims that she does not allow herself to possess when carrying out the inductive completeness argument. Put differently, the objection denies the empiricist proponent of the inductive completeness argument information about arithmetical claims that she must suppose her to possess if the objection is to get off the ground.

Acknowledgments

I am grateful to, among others, Jake Chandler, Richard Dawid, Rory Harder, Ole Hjortland, David Hyder, Norbert Gratzl, Terence Irwin, Hannes Leitgeb, Paul Rusnock, Gregory Wheeler, Timothy Williamson, and, especially, Wolfgang Spohn and Marta Sznajder for many helpful comments and suggestions on earlier versions of this paper.

My research was supported by the Canadian SSHRC through its Insight program and by the Connaught Foundation through its New Researcher program.

References

- [1] Brössel, Peter, & Eder, Anna-Maria A., & Huber, Franz (2013), Evidential Support and Instrumental Rationality. *Philosophy and Phenomenological Research* **87**, 279-300.
- [2] Buchak, Lara (2014), *Risk and Rationality*. Oxford: Oxford University Press.
- [3] Burks, Arthur W. (1963), On the Significance of Carnap's System of Inductive Logic for the Philosophy of Induction. P.A. Schilpp (ed.), *The Philosophy of Rudolf Carnap*. La Salle, IL: Open Court, 739-760.
- [4] Carnap, Rudolf (1934), *Logische Syntax der Sprache*. Vienna: Springer.
- [5] Carnap, Rudolf (1936), Testability and Meaning. *Philosophy of Science* **3**, 419-417.
- [6] Carnap, Rudolf (1937), Testability and Meaning – Continued. *Philosophy of Science* **4**, 1-40.
- [7] Carnap, Rudolf (1947), *Meaning and Necessity. A Study in Semantics and Modal Logic*. Chicago: University of Chicago Press.
- [8] Carnap, Rudolf (1950), *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- [9] Carnap, Rudolf (1952), *The Continuum of Inductive Methods*. Chicago: University of Chicago Press.

- [10] Carnap, Rudolf (1962), *Logical Foundations of Probability*. 2nd ed. Chicago: University of Chicago Press.
- [11] Carnap, Rudolf (1963a), Intellectual Autobiography. In P.A. Schilpp (ed.), *The Philosophy of Rudolf Carnap*. La Salle, IL: Open Court, 3-84.
- [12] Carnap, Rudolf (1963b), Replies and Systematic Expositions. Probability and Induction. In P.A. Schilpp (ed.), *The Philosophy of Rudolf Carnap*. La Salle, IL: Open Court, 966-998.
- [13] Carnap, Rudolf (1968), Inductive Logic and Inductive Intuition. In I. Lakatos (ed.), *The Problem of Inductive Logic*. Amsterdam: North-Holland, 258-267.
- [14] Carnap, Rudolf (manuscript), How can induction be justified? University of Pittsburgh Libraries. Special Collections Department.
- [15] Carnap, Rudolf & Hahn, Hans & Neurath, Otto (1929), *Wissenschaftliche Weltauffassung. Der Wiener Kreis*. Wien: Artur Wolf Verlag.
- [16] Gettier, Edmund L. (1963), Is Justified True Belief Knowledge? *Analysis* **23**, 121-123.
- [17] Gödel, Kurt (1930), Die Vollständigkeit der Axiome des logischen Funktionenkalküls. *Monatshefte für Mathematik und Physik* **37**, 349-360.
- [18] Gödel, Kurt (1931), Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik Physik* **38**, 173-198.
- [19] Gold, E. Mark (1967), Language Identification in the Limit. *Information and Control* **10**, 447-474.
- [20] Goodman, Nelson (1946), A Query on Confirmation. *Journal of Philosophy* **43**, 383-385.
- [21] Goodman, Nelson (1947), On Infirmities of Confirmation-Theory. *Philosophy and Phenomenological Research* **8**, 149-151.
- [22] Goodman, Nelson (1954), *Fact, Fiction, Forecast*. Cambridge, MA: Harvard University Press.

- [23] Goodman, Nelson (1983), *Fact, Fiction, Forecast*. 4th ed. Cambridge, MA: Harvard University Press.
- [24] Grice, Herbert P. (1975), Logic and Conversation. In D. Davidson & G. Harman (eds), *The Logic of Grammar*. Encino, CA: Dickenson, 64-75.
- [25] Haack, Susan (1976), The Justification of Deduction. *Mind* **85**, 112-119.
- [26] Hume, David (1739/1896), *A Treatise of Human Nature*. Ed. by L.A. Selby-Bigge. Oxford: Clarendon Press.
- [27] Hume, David (1748/1993), *An Enquiry Concerning Human Understanding*. Ed. by E. Steinberg. Indianapolis: Hackett.
- [28] Jackson, Frank (1979), On Assertion and Indicative Conditionals. *The Philosophical Review* **87**, 565-589.
- [29] Joyce, James M. (1998), A Nonpragmatic Vindication of Probabilism. *Philosophy of Science* **65**, 575-603.
- [30] Joyce, James M. (2009), Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In F. Huber & C. Schmidt-Petri (eds.), *Degrees of Belief*. Dordrecht: Springer, 263-300.
- [31] Kant, Immanuel (1902), *Kants gesammelte Schriften*. Ed. by the Königlich Preussische (now Deutsche) Akademie der Wissenschaften. Berlin: G. Reimer (now de Gruyter).
- [32] Kelly, Kevin T. (1996), *The Logic of Reliable Inquiry*. Oxford: Oxford University Press.
- [33] Kelly, Kevin T. (2007), A New Solution to the Puzzle of Simplicity. *Philosophy of Science* **74**, 561-573.
- [34] Kemeny, John G. (1963), Carnap's Theory of Probability and Induction. P.A. Schilpp (ed.), *The Philosophy of Rudolf Carnap*. La Salle, IL: Open Court, 711-738.
- [35] Kraus, Sarit & Lehmann, Daniel & Magidor, Menachem (1990), Nonmonotonic Reasoning, Preferential Models, and Cumulative Logics. *Artificial Intelligence* **40**, 167-207.

- [36] Korsgaard, Christine M. (2008), The Normativity of Instrumental Reason. In Korsgaard, Christine M. (2008), *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford: Oxford University Press, 1-41.
- [37] Lenz, John W. (1956), Carnap on Defining “Degree of Confirmation”. *Philosophy of Science* **23**, 230-236.
- [38] Lewis, David (1986), Probabilities of Conditionals and Conditional Probabilities II. *The Philosophical Review* **95**, 581-589.
- [39] Martin, Eric & Osherson, Daniel (1998), *Elements of Scientific Inquiry*. Cambridge, MA: MIT Press.
- [40] Nagel, Ernest (1963), Carnap’s Theory of Induction. P.A. Schilpp (ed.), *The Philosophy of Rudolf Carnap*. La Salle, IL: Open Court, 785-826.
- [41] Percival, Philip (2002), Epistemic Consequentialism. *Proceedings of the Aristotelian Society*, Supplementary Volume **76**, 121-151.
- [42] Putnam, Hilary (1963), Degree of Confirmation and Inductive Logic. P.A. Schilpp (ed.), *The Philosophy of Rudolf Carnap*. La Salle, IL: Open Court, 761-784.
- [43] Quine, Willard van Orman (1986), Reply to Morton White. In L. Hahn & P. Schilpp (eds.), *The Philosophy of W. V. Quine*. La Salle: Open Court.
- [44] Reichenbach, Hans (1938), *Experience and Prediction. An Analysis of the Foundations and the Structure of Knowledge*. Chicago: University of Chicago Press.
- [45] Reichenbach, Hans (1940), On the Justification of Induction. *Journal of Philosophy* **37**, 97-103.
- [46] Skyrms, Brian (2000), *Choice and Chance: An Introduction to Inductive Logic*. 4th ed. Belmont, CA: Wadsworth Thomson Learning.
- [47] Stalker, Douglas F. (ed.) (1994), *Grue! The New Riddle of Induction*. Chicago: Open Court.
- [48] Stalnaker, Robert C. (2002), Epistemic Consequentialism. *Proceedings of the Aristotelian Society*, Supplementary Volume **76**, 153-168.

- [49] Strawson, Peter F. (1952), *Introduction to Logical Theory*. London: Methuen.
- [50] Strawson, Peter F. (1958), On Justifying Induction. *Philosophical Studies* **9**, 20-21.