Time and the Old Riddle

Marius Backmann, University of Konstanz

November 26, 2018

marius.backmann@uni-konstanz.de

Contents

1 Introduction

However we want to precisely define what inductive inferences are, it is clear that any empirical inductive inference includes an inference from a sample of observed past or present instances to a population of unobserved instances, be they past, present, or future, or all of the above. In this paper, I argue that whatever stance one might take on whether a justification of induction is possible, inferences to future instances are more problematic than inferences to past of unobserved instances. This is particularly the case for any justification of induction which aims to show that induction is in any way truth-conducive, and relies on a set of rules for the sampling, such as the Law of Large Numbers-approach [LLN]. As we will see, the arguments presented here are not restricted to the LLN-view, though.

Inferences to the future are more problematic than inferences to the merely unobserved due to two issues – one general, and one that applies only to temporal ontologies with an ontologically open future. Firstly, there is the problem that the sample from which to infer information about the population is necessarily not random. Secondly, there is the issue that if one proposes a temporal ontology with an open future, it is not only ontologically open what future facts will come about, but also there is no population yet about which to inductively one could infer anything. As we will see, this problem will turn out to be a lot less trivial as it might seem at first sight.

I will begin by briefly delineating what I take inductive inferences and the Old Riddle of Induction to be. Following this, I will briefly outline the LLN-approach and argue why I take it as an instance of a whole tradition or attempts to justify induction, which all suffer from the same two problems as presented here, regardless of their differences. I will then go on to present the two main arguments, before I discuss two likely objections to the arguments presented here.

2 Inductive Inferences and the Old Riddle

For our discussion here, it is advisable to find a suitably broad conception of inductive inference so as not to predetermine the outcome of our argument with our definition of the subject matter. Inductive inferences are a subset of all ampliative inferences. An inference is ampliative if the content of its conclusion goes beyond the content of the premises.

What unites all the different forms of inductive inferences is that they involve an inference from the nature of a sample to the nature of a whole population in general, or to a particular instance of such a population.[1] All of these inferences are inferences from a set of observed instances to a set of unobserved instances. But another important dimension according to which these inferences can differ in practice that often gets muddled in the debate about induction is whether the conclusion encompasses future unobserved instances. Not all inductive inferences need to refer to future cases. Consider the following case: You are presented with a large bag filled with colourful balls. You are instructed to draw ten balls from the bag and then infer the distribution of colours of all balls in the bag. You draw ten balls, and they are all red, from which you infer that all of the balls in the bag are red. This inference is different from an inference where, e.g., you infer that the *next* ball you draw will be red. The original conclusion does not refer to any future instances, but to the state of the world *now*. Similarly, when we infer the traits of some long extinct species from observing a sample of the fossils we have observed so far, we infer information about past, not future instances. In contrast to this, there are obviously many predictive inductive inferences that go from observed past or present instances to future instances: If I observe that all past attempts of our university cantina to produce an edible ratatouille have been unsuccessful, I am justified to infer that their next attempt will be unsuccessful as well.

At first sight, you might ask in which sense there even is a difference between the cantina and the colourful balls case. In both cases, we inductively

[1]Whether these inference patterns can be reduced to one another is an interesting question, but not important in our case.

infer the state of either the distribution of balls in the bag or the state of a future ratatouille, which both have not been observed yet. So in a sense, both inferences concern the future and can only be supported of disconfirmed in the future, relative to the instant the inference is made in. The difference is that what we infer the state of has a different temporal status: the balls in the bag are all already present, whereas the next ratatouille that our university cantina will produce and inevitably slaughter is still in the future. So while the next observations that would support or disconfirm our hypotheses on either balls or ratatouille both lie in the future, the cases are different because in the balls case, the population we infer anything about is already present, whereas in the ratatouille case it isn't.

To wrap up, we can take inductive inferences to be inferences from the nature of an observed sample of past or present instances to a population of past, present, or future, or all of the above, instances, or a subset thereof. To delineate what exactly the problem of induction amounts to is again a very complex matter, with which we will not burden ourselves here. For the purpose of this paper, it suffices to say that the inductive inferences *qua* being ampliative inferences are in need of justification, since ampliative inferences are not truth-preserving as deductive inferences are. What such a justification would need to entail is not uncontroversial and depends on the way one exactly formulates the problem.[2]

Contemporary attempts to justify induction usually take one of the three forms: Firstly, there exist a number of attempts to reduce inductive inferences to inference patterns that are supposedly more reliable. Laurence BonJour for example holds that induction can be reduced to a two-step inference. First, from a set of observed instances of *Fs* that were *Gs* we infer via inference to the best explanation (IBE) that there is an according regularity that all *Fs* are *Gs*. From this regularity, we then deductively infer that all future or unobserved *Fs* are *Gs*.[3] David Armstrong has a similar proposal: from our observation that so far, all *Fs* have been *Gs*, we infer via IBE that *F-ness* necessarily brings along *G-ness*, from which necessary

[2] *redacted for review

[3] BonJour (1998)

connection we are allowed to deductively infer that all *Fs* are *Fs*.[4] A second, more popular, approach to justify induction would be to show how, without reducing induction to other types of inferences, one can show that induction is at least truth-conducive. The LLN-approach which we will discuss as an example below is such an attempt. A third possible way to try to solve the problem is to demonstrate that regardless of whether induction can be shown to be truth-conducive or not, inductive practice can at least to be shown to be a rational endeavour. One possibility to do so is to identify inductive reasoning as a type of Bayesian conditionalisation, and then applying the arguments to the rationality of conditionalisation such as the Dutch Strategies Argument or the Expected Epistemic Utility Argument to inductive reasoning.[5]

In the following, I will concentrate on the LLN-approach to justify induction. Not because I believe that it is the best possible strategy to justify induction – in fact I believe it fails like all attempts that try to show that induction is truth-conducive – but because it can serve as an example for any attempt to justify induction by showing that induction is truth-conducive without reducing it to any other inference pattern.

3 Induction and the (Open) Future

The LLN-approach to induction is an instance of a whole tradition of attempts solve the riddle of induction by applying classical probability theory. The view has been defended, for example, by David Stove and Donald Williams.[6] The basic idea is that if one draws a large enough random sample from a population, it is very likely that the population will resemble the population. In fact, if we drew every logically possible large sample from a population in order to establish the distribution of a certain trait, then in a majority of these samples, the distribution of that trait will fall into a very

[4] Armstrong (1983), 54-59.

[5] See e.g. Teller (1973) for the dutch strategies arguments, or GreavesWallace2006 for the expected epistemic utility argument.

[6] See e.g. Stove (1986) and Williams (1947).

small margin from the original distribution in the population. Given that fact, we are supposedly justified to infer that a large sample will resemble the population. However, in order for us to be justified that the sample will resemble the population, two conditions have to be fulfilled: firstly, the sample must be random. And secondly, the population must be finite. As it turns out, both conditions are problematic if the inductive inference is one from past and present to future instances. Let us now turn to both of these objections, before we discuss why this problem is pervasive for any justification of induction that aims to show that induction is truth-conducive.

These two conditions hold for any attempts to justify induction that emphasise the importance of the long run. Although, for instance, Reichenbach's and Salmon's view on induction is categorically different from the LLN-approach in that Reichenbach and Salmon both claim that although it can never be conclusively demonstrated that nature is uniform, they hold that if any mode of ampliative inference is successful, induction is.[7] Salmon proposes the following inference rule:

> [G]iven m/n of observed A are B, [...] infer that the 'long run' frequency of B among A is m/n.[8]

This inference rule, just like the inference rules of the LLN-approach, presupposes that the population is finite and the sample random. The same holds for any attempt to justify induction that holds that there is any long run resemblance between a large enough sample and the population the sample was taken from. Let us now turn to the issues with accounts like these when it comes to inductive predictions.

3.1 Randomness

For any inductive inference, it should go without saying that if a sample is skewed, and especially if we know it to be skewed, we cannot be confident enough that it is one of the majority of logically possible distributions that

[7] See Reichenbach (1935), and Salmon (1974).
[8] Salmon (1974), 50.

fall within a very small margin of the population. Consider our bag of balls again. If we know that there are balls of different colours in a bag, and if we know that the balls are not distributed evenly in that bag, but rather layered according to their colour, we are not justified to infer the proportion of balls in the bag from a sample if we e.g. only drew balls from a very small region in the top left corner.

Problematically, the same problem arises when we infer from past and present to future instances: any sample we draw in order to make a prediction is not random, but a temporally ordered and temporally restricted set of instances we observed over a certain period of time. Say the population we are trying to inductively infer anything about is indeed made up of individuals which might change over time – either individually, in the sense that the traits of any individual within the population might change – or generally, in the sense that the individuals themselves do not change, but it changes which individuals with which particular traits are present at a given time. Then a spatiotemporally restricted and ordered non-random sample will not allow us to justify inferring the characteristics of a – possibly changeable – population.

Consider an example. You take your canoe and go out paddling on your local lake, which is a popular boating spot amongst local people and tourists alike. There is a number of canoes, kayaks, paddle boards, and sail- and motor boats on the lake, and you start to wonder what the percentage of canoes amongst the water craft on the lake is on that day. The lake is too big to be overseen in its entirety, so you paddle along the lake, and while you do that, you count the number of canoes and other craft. At the end of the day, you observed a large number of boats at different spots on the lake, and of all observed boats about 15 percent were canoes. You conclude that your sample was large enough, and sampled on a large enough variety of regions of the lake, to infer that the total percentage of canoes on the entire lake is indeed close to 15 percent. This inference would be justified according to the LLN-approach – however, the inference concerns only the population of canoes present on that lake on that day, and you would not immediately infer that that will be the percentage of canoes on that lake for time eternal, just like you would not conclude that that your observation says anything about the number of canoes on all lakes, not just your local one, no matter how many boats you observed that day. Your sample, though large, was not random: It was spatially and temporally restricted, and maybe that day 2 canoe clubs where on the lake, a lot of canoeing tourists where present that day, canoes have been on sale at the local boat store 2 months, and it was a nice, sunny, calm day. And maybe all the hip crowd will turn to stand-up paddle boards in the future, foolishly abandoning their canoes.

But maybe it is hardly surprising that predictions are more problematic than non-predictive inferences: after all, things change. So what is the big deal? The problem is that the attempts to justify induction that claim that induction is truth-conducive if we only observe enough instances require the sample to be random, which a temporally ordered one can never be. If the population we draw a sample of is composed of past, present, and future instances, then obviously any temporally ordered and restricted sample will not meet the internal criteria of any such approach to induction. Just like we cannot infer the composition of balls in a bag if we only drew a sample from one corner, we are not justified to infer anything about the percentage of past, present, and future canoes amongst all boats on lake from a sample of one day. However, this entails that according to the LLN-account, a whole subsection of inductive inferences, i.e. inferences from premises referring to past and present instances to a conclusion that refers to future instances, are not straightforwardly justified. Simply observing a large number of instances will not do to justify inferring anything about other instances, if the observed ones are past and present and the inferred instances are in the future.

So far, our discussion here has made no reference to any particular notion of time. However, things get even more complicated if we add temporal ontology to the mix. As we will see below, temporal ontologies with an open future result in an additional problem regarding the population from which we draw our sample and about which we make our predictions.

3.2 The Open Future

Different temporal ontologies warrant different notions on the openness of the future. Let us, with a very broad brush, distinguish between three main rival temporal ontologies: Eternalism, Presentism, and the Growing Block Theory (GBT). In Eternalism, all events, past, present, and future, are all equally real. In B-theoretic eternalism (and thus excluding the Moving Spotlight Theory), there is no objective fact of the matter as to which facts are past, present, or future: facts are temporally ordered, but there is no objective progressing present. In Presentism, only the present facts are real, whereas the past ones are not real anymore and the future ones are not real yet. Typically, Presentism is A-theoretic and includes a notion of objective and progressing presentness. The GBT is a combination of the two views: with Eternalism, the GBT shares the characteristic that there is a block of past and present facts, which are all equally real. The future facts, however, are not real yet: they come into existence as the objective present progresses, which accounts for a growing of the block of facts. *Prima facie*, Presentism and the GBT have a more robust notion of the openness of the future: while in Eternalism, all facts are already in place, even the ones that are in the future from an arbitrary temporal location in the block, the same cannot be said about the GBT and Presentism. There, the future is ontologically open in the sense that the future facts are not fixed because they are not real yet.

This complicates any solution to the problem of induction like the ones discussed here greatly, or any approach that emphasises the long run like Reichenbach and Salmon. Say all the ravens we have observed so far have been black. If from this, we wanted to infer that *all* ravens are black, what would "all ravens" refer to? All existing ravens? If that is the case, we should discount the future ravens in GBT and Presentism, because there are no future ravens – not yet at least. So if we wanted to make an inductive prediction about the future based on a large sample from a population, what do we take the population to be? If the population is supposed to include future instances, then we are in the curious position that according to temporal ontologies with an open future, the population is composed partly of real ravens, and partly of ones that aren't real yet. In a standard reading of these temporal ontologies, there is no fact of the matter *now* about the colour of future ravens, because there are no real future ravens. Hence, there is no fact of the matter now whether future ravens will be black, yellow, green, or pink.

Consider our bag full of balls again. There are 20.000 balls in the bag, and you are asked to draw 3000 of those, all of which turn out red. According to the LLN-approach, you are justified to infer that most likely, all balls in the bag are red, provided the sample was random. But imagine now that after this inference, 300.000 new balls are added to the bag, step by step, or ball by ball. Now, the population of balls in the bag has changed, and keeps changeing. So the population whose proportion of red balls you are trying to infer, that is the original 20.000 balls, is a different one than the population of balls in the bag after we added 300.000 more balls. Clearly, even given the large sample size, we are not justified in inferring anything about a different population than the one we drew our sample from. For all we know, all 300.000 new balls might be blue.

There are at least two possible objections to this. One possible way to deal with the problems raised in this paper would be to find a way to restrict the range of possible future facts that will occur and be added to the block of real facts. Such a proposal has been made by R.A. Briggs and Graeme A. Forbes. Briggs and Forbes claim that even in the GBT, there is a fact of the matter *now* about what will be the case, even if there are no future facts. They claim that if we proposed any sort of necessitarianism, the necessary connections that are instantiated *now* restrict the possibilities of what will or could come about. And these necessary connections and the set of real past or present facts together could serve as a truth-maker for future tensed assertions, given that determinism is true (at least for the relevant facts)[9] Without going into too much detail, this solution is problematic. How do we know, e.g., that the necessary connections that are instantiated now will

[9]Briggs and Forbes (2012)

not change in the future? If Fs and Gs are so related now that any F necessarily brings about a G, how do we know that this will also be the case in the future? And if it is possible that necessary connections could change, then the necessary connections instantiated now cannot restrict the range of possible future facts. Helen Beebee, for example, has invoked such an argument against necessitarian attempts such as David Armstrong's to justify induction.[10] However that discussion turn out, it seems that people like Briggs and Forbes, who try to invoke necessary connections to restrict the range of possible future facts, also have to argue for the immutability of necessary connections.

Another possible objection would be to invoke some version of Ersatzism. Routinely, proponents of temporal ontologies with an open future (or unreal past) are faced with some variant of a truth-maker objection: If all propositions require truth-makers to have a truth value, then future contingents (or, in the case of presentism, even past tense assertions), have no truth value, because there are no future (or past) facts that could act as a truth-maker. A prominent answer to this challenge is to adopt a form or Ersatzism:[11] in temporal ontologies with an open future, the proposition "the next raven we will observe will be black", cannot be made true at the time of utterance by a real future fact that at some future time t, we observe a black raven. In Ersatzism, the role of truth-maker of that assertion is taken over by a timeless fact that we will observe this particular black raven at some abstract, "ersatz" future time. To put it simply, there might not be a future time t where it is the case that we observe a black raven, but it is a fact now that there will be a time where we observe that raven. A natural response to our problem of the growing population in the case of predictive induction would be to say that while the population about which we infer anything might not exist yet, it is the case *now* that there will be a population with a certain set of characteristics. Ersatz times thus could provide us with an ersatz population.

I will take no stance here on whether Ersatzism is a satisfactory answer

[10]Beebee (2011)

[11]See Bourne (2006) and Crisp (2007).

to the truth-maker challenge to temporal ontologies with an open future ? although I am highly sceptical about the prospects of such a view. In our case, Ersatzism would have to do a lot of heavy ontological and epistemic lifting. Ontologically, Ersatzism would have to provide us with an ersatz population, existing at ersatz times, which could not only serves as truth-makers for our predictions, but also serve as an ersatz population from which to draw a sample in the first place. Moreover, remember that the population has to be finite for any long run justification of induction to work. So if there is no future population, then there is no no sense in which we could even reasonably ask whether it is finite. If there are no real unicorns, the question of whether there is a finite or an infinite number of real unicorns does not even make sense. If the future is genuinely ontologically open, it is *ontologically* indeterminate whether the population is finite or not. It is thus ontologically indeterminate whether our inferences pertaining to future instances are justifi*able* at all. Ersatzism would have to provide us with a robust sense that there is a finite ersatz population of future ravens at ersatz future times that could help us justify inductive predictions. But even if Ersatzism is a satisfactory ontological answer to these worries, the epistemic problem still remeins: even if a finite future ersatz population could make it the case that now, our predictive inductive inferences are justifiable, we do not *know* whether the ersatz population is finite or not ? whether there will be finitely many ravens seems to be the sort of thing one cannot know *a priori*, but has to infer inductively. This would result in a vicious circle: to know whether our future ersatz population is finite or not, and hence, to know whether inductive predictions are justifiable at all, we would need to make an inductive prediction.

It seems like the discussed approaches to the Old Riddle do not merely fail for predictive inductive inferences, but that they are not even straight-forwardly applicable, if they treat the population about which we make inductive predictions as real, or fixed, at the moment the inductive prediction is made.

4 Conclusion

So far, we have only discussed approaches to solve the problem of induction that emphasise the long run and claim that a sufficiently large, random sample will likely resemble the population the sample was drawn from. However, the problems raised here pertain to any attempt to justify induction that try to demonstrate that induction is truth-conducive. As has been a staple in the debate ever since Hume's sceptical argument, we need a notion of uniformity of nature for any such attempt to justify induction: if we are to be justified to infer from an observed sample to a larger population, we need the notion that this larger population does not change randomly. If there is a block of past, present, or future facts, there is at least a hope to get around the objection that we don't know whether the regularities will remain stable. If we were to find a response to the argument that our sample is not random of we infer from a temporally ordered and temporally restricted set of instances to a population containing future ones, then the LLN-approach, e.g., suitably amended, might be able to solve that problem. If, however, there are no future facts, and if there is no fact of the matter now that restricts what will happen, then there is no way to secure that the future will resemble the present and the past. Thus, a temporal ontology with open future renders any attempt to solve the problem of induction that invokes some notion of the uniformity of nature unsuccessful.

References

David M. Armstrong. *What is a Law of Nature?* Cambridge University Press, Cambridge, 1983.

Helen Beebee. Necessary connections and the problem of induction. *Noûs*, 45(3):504–527, 2011.

Laurence BonJour. *In Defense of Pure Reason. A Rationalist Account of A Priori Justification.* Cambridge Studies in Philosophy. Cambridge University Press, Cambridge, 1998.

Craig Bourne. A theory of presentism. *Canadian Journal of Philosophy*, 36 (1):1–23, 2006.

Rachel Briggs and Graeme A. Forbes. The real truth about the unreal future. In Karen Bennett and Dean W. Zimmerman, editors, *Oxford Studies in Metaphysics*, volume 7, pages 257–304. Oxford University Press, Oxford, 2012.

Thomas M. Crisp. Presentism and the groundng objection. *Noûs*, 41(1): 90–109, 2007.

Hans Reichenbach. *Wahrscheinlichkeitslehre – Eine Untersuchung über die logischen und mathematischen Grundlagen der Wahrscheinlichkeitsrechnung.* A.W. Sijthoff's Uitgevermaatschapij N.V., Leiden, 1935.

Wesley Salmon. The concept of inductive inference. In Richard Swinburne, editor, *The Justification of Induction*, pages 49–57. Oxford University Press, Oxford, 1974.

David Stove. *The Rationality of Induction.* Oxford University Press, Oxford/New York, 1986.

Paul Teller. Conditionalization and observation. *Synthese*, 26:218–258, 1973.

Donald Williams. *The Ground of Induction.* Harvard University Press, Cambridge, MA, 1947.